arXiv:1602.06260v1 [cs.CC] 19 Feb 2016

Checking Dynamic Consistency of Conditional Hyper Temporal Networks via Mean Payoff Games

Hardness and (pseudo) Singly-Exponential Time Algorithm

Carlo Comin* Romeo Rizzi†

Abstract

Conditional Simple Temporal Network (CSTN) is a constraint-based graph-formalism for conditional temporal planning. It offers a more flexible formalism than the equivalent CSTP model of Tsamardinos, et al., from which it was derived mainly as a sound formalization. Three notions of consistency arise for CSTNs and CSTPs: weak, strong, and dynamic. Dynamic-consistency is the most interesting notion, but it is also the most challenging and it was conjectured to be hard to assess. Tsamardinos, et al. gave a doubly-exponential time algorithm for checking the dynamic-consistency of CSTNs and to produce, in the positive case, an exponential sized dynamic execution strategy. The CSTN model may be viewed as an extension of Simple Temporal Networks (STN) (Dechter, et al.), i.e., directed weighted graphs where nodes represent events to be scheduled in time and arcs represent temporal distance constraints between pairs of events. Recently, STNs have been generalized into *Hyper Temporal Networks* (HyTNs), by considering weighted directed hypergraphs, where each hyperarc models a *disjunctive* temporal constraint (*hyper-constraint*); where, being directed, the hyperarcs can be either *multi-head* or *multi-tail*. The computational equivalence between checking the consistency of HyTNs and determining the winning regions in Mean Payoff Games (MPGs) was pointed out as well; the latter is a family of two-player infinite games played on finite graphs, well known for having applications in model-checking and formal verification. In this work we introduce the *Conditional Hyper Temporal Network (CHyTN)* model, which is a natural extension and generalization of both the CSTN and the HyTN model.Our contribution goes as follows. We start by proving that deciding whether a given CSTN or CHyTN is dynamically consistent is coNP-hard. Then, we offer a proof that deciding whether a given CHyTN is dynamically consistent is PSPACE-hard, provided that the input instances are allowed to include both multi-head and multi-tail hyperarcs. In light of this, we continue our study by focusing on CHyTNs that allow only multi-head or only multi-tail hyperarcs, and we offer the first deterministic (pseudo) singly-exponential time algorithm for the problem of checking the dynamic-consistency of such CHyTNs, also producing a dynamic execution strategy whenever the input CHyTN is dynamically consistent. Since CSTNs are a special case of CHyTNs, this provides as a byproduct the first sound-and-complete (pseudo) singly-exponential time algorithm for checking dynamic-consistency in CSTNs. The proposed algorithm is based on a novel connection between CSTNs/CHyTNs and Mean Payoff Games. The presentation of the connection between CSTNs/CHyTNs and MPGs is mediated by the HyTN model. Finally, in order to analyze the algorithm, we introduce a refined notion of dynamic-consistency, named ϵ-dynamic-consistency, and present a sharp lower bounding analysis on the critical value of the reaction time $\hat{\varepsilon}$ where a CSTN/CHyTN transits from being, to not being, dynamically consistent. The proof technique introduced in this analysis of $\hat{\varepsilon}$ is applicable more generally when dealing with linear difference constraints which include strict inequalities.

*Department of Mathematics, University of Trento, Trento, Italy; LIGM, Université Paris-Est, Marne-la-Vallée, Paris, France. E-mail: carlo.comin@unitn.it

†Department of Computer Science, University of Verona, Verona, Italy. E-mail: romeo.rizzi@univr.it

Keywords: Conditional Simple Temporal Networks, Conditional Hyper Temporal Networks, Dynamic Consistency, Mean Payoff Games, STNs, Hyper Temporal Networks, Singly-Exponential Time, Reaction Time.

1 Introduction and Motivation

In many areas of Artificial Intelligence (AI), including temporal planning and temporal scheduling, the representation and the management of quantitative temporal aspects is of crucial importance (see e.g., [25, 26, 17, 3, 10, 9]). Examples of possible quantitative temporal aspects include: constraints on the earliest start time and latest end time of activities, constraints over the minimum and maximum temporal distance between activities, etc. In many cases these constraints can be represented as instances of *Simple Temporal Networks* (STNs) [16], i.e., directed weighted graphs, where nodes represent events to be scheduled in time and arcs represent temporal distance constraints between pairs of events. Recently, STNs have been generalized into *Hyper Temporal Networks* (HyTNs) [11, 12], which is a strict generalization of STNs, introduced to overcome the limitation of considering only conjunctions of constraints, but maintaining a practical efficiency in the consistency check of the instances. In a HyTN a single temporal hyperarc constraint may be defined as a set of two or more maximum delay constraints which is satisfied when at least one of these delay constraints is satisfied. HyTNs are meant as a light generalization of STNs offering an interesting compromise. On one side, there exist practical pseudo-polynomial time algorithms for checking the consistency and computing feasible schedules for HyTNs. On the other side, HyTNs offer a more powerful model accommodating natural constraints that cannot be expressed by STNs like "*Trigger off exactly δ min before (after) the occurrence of the first (last) event in a set.*" To give more details, HyTNs are weighted directed hypergraphs, where each hyperarc models a *disjunctive* temporal constraint called *hyper-constraint*. The computational equivalence between checking the consistency of HyTNs and determining the winning regions in *Mean Payoff Games* (MPGs) [18, 29, 5, 14] was also pointed out in [11, 12], where the approach was shown to be robust by extensive experimental evaluations [11, 12, 4]. MPGs are a family of two-player infinite games played on finite graphs, well known for having theoretical interest in computational complexity, being it one of the few (natural) problems lying in $\mathsf{NP} \cap \mathsf{coNP}$, as well as various applications in model-checking and formal verification [20].

However, another aspect of crucial importance in the representation's realm of quantitative temporal aspects of systems is the fact that *conditional* temporal constraints pose a serious challenge for *conditional* temporal planning, because it is necessary for a conditional planner to determine whether a candidate plan will satisfy the specified temporal constraints. This can be difficult, as in fact the temporal assignments that satisfy the constraints associated with one conditional branch may fail to satisfy the constraints along a different branch, see e.g., [28]. The present work unveils that HyTNs and MPGs are a natural underlying combinatorial model for checking the consistency of certain conditional temporal problems that are known in the literature and also useful in some practical applications of temporal planning, especially for managing the temporal aspects of Workflow Management Systems (WfMSs) [3, 10] and for modeling Healthcare's Clinical Pathways [9]. So, we shall begin our study by focusing on *Conditional Simple Temporal Problems* (CSTP) [28] and on their graph-based counterpart *Conditional Simple Temporal Networks* (CSTN) [21], a suitable constraint-based model for conditional temporal planning. The CSTN formalism extends STNs in that: (1) some of the nodes are called *observation events* and to each of them is associated a boolean variable, to be disclosed only at execution time; (2) *labels* (i.e. conjunctions over the literals) are attached to all nodes *and* constraints, to indicate the situations in which each of them is required. The planning agent must schedule all the required nodes, meanwhile respecting all the required temporal constraints among them. This extended framework allows for the off-line construction of conditional plans that are guaranteed to satisfy complex temporal constraints. Importantly, this can be achieved even while allowing for the decisions about the

precise timing of actions to be postponed until the execution time, in a least-commitment manner, thereby adding flexibility and making it possible to adapt the plan dynamically, during execution, in response to the observations made, see e.g., [28] for further details and examples.

Three notions of consistency arise for CSTNs: weak, strong, and *dynamic*. Dynamic-consistency (DC) is in fact the most interesting one, as it requires the existence of conditional plans where decisions about the precise timing of actions are postponed until execution time, but it anyhow guarantees that all the relevant constraints will be ultimately satisfied. Still, it is the most challenging and it was conjectured to be hard to assess by [28]. Indeed, the best algorithm which is currently known for deciding whether a given CSTN is dynamically-consistent is doubly-exponential time [28]. It first builds an equivalent Disjunctive Temporal Problem (DTP) of size exponential in the input CSTN, and then applies to it an exponential time DTP's algorithm to check its consistency. However, this approach turns out to be quite limitative in practice: to the best of our knowledge, some experimental studies have already shown that the resolution procedures, as well as the currently known heuristics, for solving general DTPs becomes quite burdensome with $\sim 30, 35$ DTP's variables, see e.g., [27, 23, 24], thus dampening the practical applicability of the DTP-based approach.

1.1 Contribution

In this work we introduce and study the *Conditional Hyper Temporal Network (CHyTN)*, which is a natural extension and generalization of both the CSTN and the HyTN model, and it is obtained by "melting" them together. One of the motivations for studying such a generalization is that to transpose to the context of *conditional* temporal planning the benefits, and opportunities for application, that arose from the introduction of the HyTN model (see e.g., [11, 12], for some motivating examples). But even more pragmatically, the main and perhaps most important contribution, in doing so, is that of being able to offer the first known sound-and-complete deterministic (pseudo) singly-exponential time algorithm for checking the dynamic-consistency of CSTNs. In summary, our contributions are as follows. After having introduced the CHyTN model formally, we start by showing that deciding whether a given CSTN or CHyTN is dynamically consistent is coNP-hard. Then, we offer a proof that deciding whether a given CHyTN is dynamically consistent is a PSPACE-hard problem, provided that the input CHyTN instances are allowed to include both multi-head and multi-tail hyperarcs. In light of this observation, we continue our study by focusing on CHyTNs that allow either only multi-head or only multi-tail hyperarcs. Concerning those kinds of CHyTNs, perhaps most importantly, we unveil a connection between the problem of checking their dynamic-consistency and that of determining the winning regions in MPGs (of a singly-exponential size in the number of propositional variables of the input CHyTN), thus providing the first sound-and-complete (pseudo) singly-exponential time algorithm for this same task of deciding the dynamic-consistency and yielding a dynamic execution strategy for CHyTNs. The resulting worst-case time complexity of our DC-Checking procedure is actually $O(2^{3|P|}|\mathcal{A}|^2|V| + 2^{4|P|}|\mathcal{A}||V|^2|P| + 2^{5|P|}|V|^3|P|)W$, where $|P|$ is the number of propositional variables, $|V|$ is the number of event nodes, $|\mathcal{A}|$ is the size (i.e., roughly, the number of arcs), and W is the maximum absolute value of the weights of the input CHyTN. The algorithm is still based on representing a given CHyTN instance on an exponential sized network, as first suggested in [28]. The difference, however, is that we propose to map CSTNs and CHyTNs on (exponential sized) HyTNs/MPGs rather than on DTPs. This makes quite a relevant difference, because the consistency check for HyTNs can be reduced to MPGs determination, as shown in [11, 12], which is amenable to practical and effective pseudo-polynomial time algorithms (indeed, in some cases the resolution methods for determining MPGs exhibit even a strongly-polynomial time behaviour, see e.g., [12, 4, 1, 6]). To summarize, we obtain an improved upper bound on the theoretical time complexity of the DC-checking of CSTNs (i.e., from 2-EXP to pseudo-$E \cap$ NE $\cap$ coNE) together with a faster DC-checking procedure, which can be used on CSTNs/CHyTNs with a larger number of propositional variables and event nodes than before. At the heart of the algorithm a suitable reduction to MPGs is mediated by the HyTN model, i.e., the algorithm decides whether a CSTN/CHyTNs is dynamically-consistent by

solving a carefully constructed MPG. As a final contribution, in order to analyze the algorithm, we introduce a novel and refined notion of dynamic-consistency, named ϵ-dynamic-consistency (where $\epsilon \in \mathbf{R}_{>0}$), and present a sharp lower bounding analysis on the critical value of the *reaction time* $\hat{\varepsilon}$ where a CSTN/CHyTNs transits from being, to not being, dynamically-consistent. We believe that this contributes to clarifying (with respect to previous literature, e.g., [28, 21]) the role played by the reaction time $\hat{\varepsilon}$ in checking the dynamic-consistency of CSTNs. Furthermore, the proof technique introduced in this analysis of $\hat{\varepsilon}$ is applicable more generally when dealing with linear difference constraints which include strict inequalities, therefore, it may be useful in the analysis of other models of temporal constraints.

A preliminary version of this article appeared in the proceedings of the TIME symposium [13]. Here, the presentation is extended as follows: (1) the definition of CSTN has been extended and generalised to that of CHyTN, in an explicit way, in order to allow for the presence of two kinds of hyperarcs as labeled temporal constraints within the input instances; (2) some further issues and pertinent properties about CSTN/CHyTNs have been observed and proved; (3) for instance, we offer a novel PSPACE-hardness result: deciding whether a given CHyTN is dynamically consistent is *hard* for PSPACE (the reduction goes from 3-CNF-TQBF), provided that the input instances are allowed to include both multi-head and multi-tail hyperarcs; (4) the proposed (pseudo) singly-exponential time algorithm is presented here in its full generality, with respect to the CHyTN model; (5) several proofs have been polished, expanded and clarified (e.g., those concerning the reaction time analysis of $\hat{\epsilon}$).

1.2 Organization

In Section 2 we recall the basic formalism, terminology and known results on STNs and HyTNs. In particular: Subsection 2.1 deals with STNs, Subsection 2.2 deals with HyTNs, and its computational equivalence with MPGs and the related algorithmic results. Section 3 surveys CSTNs and then it introduces CHyTNs, also presenting some basic properties of the model. Section 4 tackles on the algorithmics of dynamic-consistency: firstly, we provide a coNP-hardness lower bound, then, we offer a novel PSPACE-hardness lower bound. Next, the section describes the connection with HyTNs/MPGs and present a novel (pseudo) singly-exponential time DC-checking procedure. Section 5 is devoted at offering a sharp lower bounding analysis on the critical value of the *reaction time* $\hat{\varepsilon}$ where the CSTN transits from being, to not being, dynamically-consistent. In Section 6, some related works are discussed.

The paper concludes in Section 7.

2 Background

2.1 Simple Temporal Networks

In this first subsection, we introduce some definitions, notations and well-know results about graphs and conservative graphs; also, we recall the relation between the consistency property of STNs and the conservative property of weighted graphs. To start with, our graphs are directed and weighted on the arcs. Thus, if $G = (V, A)$ is a graph, then every arc $a \in A$ is a triplet (t_a, h_a, w_a), where: $t_a \in V$ is the *tail* of a, $h_a \in V$ is the *head* of a, and $w_a \in \mathbf{Z}$ is the *weight* of a, which we shall assume to be an integer number (unless otherwise specified). Moreover, since we use graphs to represent distance constraints, they do not need to have either loops (unary constraints are meaningless) or parallel arcs (two parallel constraints represent two different distance constraints between the same pair of nodes: only the most restrictive is meaningful). We also use the notations $h(a)$ for h_a, $t(a)$ for t_a, and $w(a)$ or $w(t_a, h_a)$ for w_a, when it helps. The *order* and *size* of a graph $G = (V, A)$ are denoted by $n \triangleq |V|$ and $m \triangleq |A|$, respectively. The size is actually a measure for the encoding length of G.

A *cycle* of G is a set of arcs $C \subseteq A$ cyclically sequenced as $a_0, a_1, \ldots, a_{\ell-1}$ so that $h(a_i) = t(a_j)$ if and only if $j = (i+1) \mod \ell$; it is called a *negative cycle* if $w(C) < 0$, where $w(C)$ stands for $\sum_{e \in C} w_e$. A graph is called *conservative* when it contains no negative cycle. A *potential* is a function $p : V \mapsto \mathbf{R}$. The *reduced weight* of an arc $a = (u, v, w_a)$ with respect to a potential p is defined as $w_a^p \triangleq w_a - p_v + p_u$. A potential p of $G = (V, A)$ is called *feasible* if $w_a^p \geq 0$ for every $a \in A$. Notice that, for any cycle C, $w^p(C) = w(C)$. Therefore, the existence of a feasible potential implies that the graph is conservative as $w(C) = w^p(C) \geq 0$ for every cycle C.

The Bellman-Ford algorithm [15] can be used to produce in $O(mn)$ time:

– either a proof that G is conservative in the form of a feasible potential function;

– or a proof that G is not conservative in the form of a negative cycle C in G.

When the graph is conservative, the smallest weight of a walk between the nodes is well defined, and, fixed a root node r in G, the potentials returned by the Bellman-Ford algorithm are, for each node v, the smallest weight of a walk from r to v. Moreover, if all the arc weights are integral, then these potentials are integral as well. Therefore, the Bellman-Ford algorithm provides a proof to the following theorem.

Theorem 1 ([2, 19, 15])**.** *A graph admits a feasible potential if and only if it is conservative. When all the arc weights are integral, the feasible potential is an integral function.*

An STN can be viewed as a weighted directed simple graph whose nodes are timepoints that must be placed on the real line and whose arcs express mutual constraints on the allocations of their end-points. An STN $G = (V, A)$ is called *consistent* if it admits a *feasible scheduling*, i.e., a scheduling $s : V \mapsto \mathbf{R}$ such that:

$$s(v) \leq s(u) + w(u, v), \quad \forall \text{ arc } (u, v) \text{ of } G.$$

Corollary 1 ([2, 16, 15])**.** *An STN G is consistent if and only if G is conservative.*

In this paper, we also deal with directed weighted *hypergraphs*.

Definition 1 (General-Hypergraph)**.** *A* general-hypergraph *$\mathcal{H}$ is a pair $(V, \mathcal{A})$, where V is the set of nodes, and $\mathcal{A}$ is the set of* hyperarcs*. Each hyperarc $A \in \mathcal{A}$ is either a* multi-head *or a* multi-tail *hyperarc.*

– A multi-head *hyperarc $A = (t_A, H_A, w_A)$ has a distinguished node t_A, called the* tail *of A, and a nonempty set $H_A \subseteq V \setminus \{t_A\}$ containing the* heads *of A; to each head $v \in H_A$ is associated a* weight *$w_A(v) \in \mathbf{Z}$, which is an integer number (unless otherwise specified). Fig. 1a depicts a possible representation of a multi-head hyperarc: the tail is connected to each head by a dashed arc labeled by the name of the hyperarc and the weight associated to the considered head.*

– A multi-tail hyperarc $A = (T_A, h_A, w_A)$ has a distinguished node h_A, called the head *of A, and a nonempty set $T_A \subseteq V \setminus \{h_A\}$ containing the* tails *of A; to each tail $v \in T_A$ is associated a* weight *$w_A(v) \in \mathbf{Z}$, which is an integer number (unless otherwise specified). Fig. 1b depicts a possible representation of a multi-tail hyperarc: the head is connected to each tail by a dotted arc labeled by the name of the hyperarc and the weight associated to the tail.*

The *cardinality* of a hyperarc $A \in \mathcal{A}$ is given by $|A| \triangleq |H_A \cup \{t_A\}|$ if A is multi-head, and $|A| \triangleq |T_A \cup \{h_A\}|$ if A is multi-tail; if $|A| = 2$, then $A = (u, v, w)$ is a standard arc. The *order* and *size* of a general-hypergraph $(V, \mathcal{A})$ are denoted by $n \triangleq |V|$ and $m \triangleq \sum_{A \in \mathcal{A}} |A|$, respectively.

2.2 Hyper Temporal Networks

This subsection surveys the *Hyper Temporal Network* (HyTN) model, which is a strict generalization of STNs introduced to partially overcome the limitation of allowing only conjunctions of constraints. HyTNs have been introduced in [11, 12], and the reader is referred there for an in-depth treatise of the subject. Compared

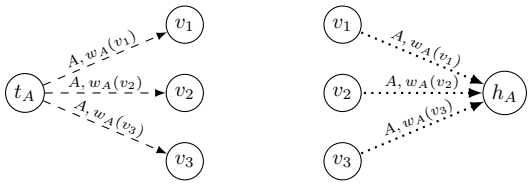

(a) Multi-Head Hyperarc $A = (t_A, H_A, w_A)$.

(b) Multi-Tail Hyperarc $A = (T_A, h_A, w_A)$.

Figure 1: A graphical representation of the two kinds of hyperarcs.

to STN distance graphs, which they naturally extend, HyTNs allow for a greater flexibility in the definition of the temporal constraints.

A HyTN is a directed weighted general-hypergraph $\mathcal{H} = (V, \mathcal{A})$ where a node represents a time point variable (timepoint), and a multi-head/multi-tail hyperarc represents a set of temporal distance constraints between the tail/head and the heads/tails, respectively. In general, we say that a hyperarc is *satisfied* when at least one of its distance constraints is satisfied. Then, we say that a HyTN is *consistent* when it is possible to assign a value to each time-point variable so that all of its hyperarcs are satisfied.

More formally, in the HyTN framework the consistency problem is defined as the following decision problem.

Definition 2 (GENERAL-HYTN-CONSISTENCY)**.** *Given a HyTN $\mathcal{H} = (V, \mathcal{A})$, decide whether there exists a scheduling $s : V \to \mathbf{R}$ such that, for every hyperarc $A \in \mathcal{A}$, the following hold:*

- *if $A = (t, h, w)$ is a standard arc, then:*

$$s(h) - s(t) \leq w;$$

- *if $A = (t_A, H_A, w_A)$ is a multi-head hyperarc, then:*

$$s(t_A) \geq \min_{v \in H_A} \{s(v) - w_A(v)\};$$

- *if $A = (T_A, h_A, w_A)$ is a multi-tail hyperarc, then:*

$$s(h_A) \leq \max_{v \in T_A} \{s(v) + w_A(v)\}.$$

Any such scheduling $s : V \to \mathbf{R}$ is called *feasible*. A HyTN that admits at least one feasible scheduling is called *consistent*.

Comparing the consistency of HyTNs with the consistency of STNs, the most important aspect of novelty is that, while in a distance graph of a STN each arc represents a distance constraint and all such constraints have to be satisfied by a feasible schedule, in a HyTN each hyperarc represents a set of one or more distance constraints and a feasible scheduling has to satisfy at least one such distance constraints for each hyperarc.

Let us survey some interesting properties about the consistency problem for HyTNs.

The first interesting property is that any integral-weighted HyTN admits an integral feasible schedule when it is consistent, as stated in the following lemma.

Lemma 1 ([12])**.** *Let $\mathcal{H} = (V, \mathcal{A})$ be an integral-weighted and consistent $HyTN$. Then $\mathcal{H}$ admits an integral feasible scheduling $s : V \to \{-T, -T+1, \ldots, T-1, T\}$, where $T = \sum_{A \in \mathcal{A}, v \in V} |w_A(v)|$.*

The following theorem states that GENERAL-HYTN-CONSISTENCY is NP-complete, in a strong sense.

Theorem 2 ([12]). GENERAL-HYTN-CONSISTENCY *is an* NP*-complete problem even if input instances* $\mathcal{H} = (V, \mathcal{A})$ *are restricted to satisfy* $w_A(\cdot) \in \{-1, 0, 1\}$ *and* $|H_A|, |T_A| \leq 2$ *for every* $A \in \mathcal{A}$*.*

Theorem 2 motivates the study of consistency problems on hypergraphs having either only multi-head or only multi-tail hyperarcs. In the former case, the consistency problem is called HyTN-Consistency, while in the latter it is called TAIL-HYTN-CONSISTENCY. In the following theorem we observe that the two problems are inter-reducible, i.e., we can solve consistency for any one of the two models in $f(m, n, W)$ time whenever we have a $f(m, n, W)$ time procedure for solving consistency for the other one.

Theorem 3 ([12]). *HyTN-Consistency and* TAIL-HYTN-CONSISTENCY *are inter-reducible by means of* $\log$*-space, linear-time, local-replacement reductions.*

In the remainder of this work we shall adopt mainly the multi-head hypergraph as our reference model; still, we will consider general-hypergraphs again in the forthcoming sections, when proving PSPACE-hardness. Let's say that, when considering hypergraphs and HyTNs, we will be implicitly referring to the multi-head hyperarcs, unless it is explicitly specified otherwise. Notably, we shall consider the following specialized notion of consistency for HyTNs.

Definition 3 (HyTN-Consistency). *Given a (multi-head) HyTN* $\mathcal{H} = (V, \mathcal{A})$*, decide whether there exists a scheduling* $s : V \to \mathbf{R}$ *such that:*

$$s(t_A) \geq \min_{v \in H_A} \{s(v) - w_A(v)\}, \quad \forall A \in \mathcal{A}. \tag{1}$$

Remark 1. *Notice that this notion of consistency for HyTNs is a strict generalization of the STN's one. In general, the feasible schedules of an STN are the solutions of a linear system and, therefore, they form a convex polytope. Since an STN may be viewed as a HyTN, the space of feasible schedules of an STN can always be described as the space of feasible schedules of a HyTN. The converse is not true because feasible schedules for a HyTN do not form a convex polytope. Let us consider, for example, a HyTN of just three nodes* x_1*,* x_2*,* x_3 *and one single hyperarc with heads* $\{x_1, x_2\}$ *and tail* x_3 *expressing the constraint* $x_3 \geq \min\{x_1, x_2\}$*;* $(0, 2, 2)$ *and* $(-2, 0, 2)$ *are both admissible schedules, but* $(1, 1, 0) = \frac{1}{2}(0, 2, 2) - \frac{1}{2}(-2, 0, 2)$ *is not an admissible schedule. In conclusion, the STN model is a special case of the Linear Programming paradigm, whereas the HyTN model is not.*

In the rest of this section, we extend the characterization of STN consistency recalled in Subsection 2.1 to HyTNs.

Definition 4 (Reduced Slack Value $w_A^p(v)$). *With reference to a potential* $p : V \to \mathbf{R}$*, we define, for every arc* $A \in \mathcal{A}$ *and every* $v \in H_A$*, the reduced slack value* $w_A^p(v)$ *as* $w_A(v) + p(t_A) - p(v)$ *and the reduced slack* w_A^p *as follows:*

$$w_A^p \triangleq \max\{w_A^p(v) \mid v \in H_A\}.$$

A potential p *is said to be* feasible *if and only if* $w_A^p \geq 0$ *for every* $A \in \mathcal{A}$*.*

Again, as it was the case for STNs, a mapping $\phi : V \to \mathbf{R}$ is a feasible potential if and only if it is a feasible schedule. In order to better characterize feasible schedules, we introduce a (generalised) notion of *negative cycle*.

Definition 5 (Negative Cycle). *Given a multi-head HyTN* $\mathcal{H} = (V, \mathcal{A})$*, a cycle is a pair* $(S, \mathcal{C})$ *with* $S \subseteq V$ *and* $\mathcal{C} \subseteq \mathcal{A}$ *such that:*

1. $S = \bigcup_{A \in \mathcal{C}} (H_A \cup \{t_A\})$ *and* $S \neq \emptyset$*;*

2. $\forall v \in S$ *there exists an unique* $A \in \mathcal{C}$ *such that* $t_A = v$.

Moreover, we let $a(v)$ *denote the unique arc* $A \in \mathcal{C}$ *with* $t_A = v$ *as required in previous item 2. Every infinite path in a cycle* $(S, \mathcal{C})$ *contains, at least, one finite cyclic sequence* $v_i, v_{i+1}, \ldots, v_{i+p}$*, where* $v_{i+p} = v_i$ *is the only repeated node in the sequence. A cycle* $(S, \mathcal{C})$ *is negative if the following holds:*

$$\sum_{t=1}^{p-1} w_{a(v_t)}(v_{t+1}) < 0, \text{ for any finite cyclic sequence } v_1, v_2, \ldots, v_p.$$

There are two results about negative cycles as stated in the following lemmas.

Lemma 2 ([12])**.** *A HyTN with a negative cycle admits no feasible schedule.*

At first sight, it may appear that checking whether $(S, \mathcal{C})$ is a negative cycle might take exponential time since one should check a possibly exponential number of cyclic sequences. The next lemma shows instead that it is possible to check the presence of a negative cycle in polynomial time.

Lemma 3 ([12])**.** *Let* $(S, \mathcal{C})$ *be a cycle in a HyTN. Then checking whether* $(S, \mathcal{C})$ *is a negative cycle can be done in polynomial time.*

A hypergraph $\mathcal{H}$ is called *conservative* when it contains no negative cycle.

In the next paragraphs we will recall the existence of a pseudo-polynomial time algorithm that always returns either a feasible scheduling or a negative cycle, thus extending the validity of the classical good-characterization of STN consistency to general HyTN consistency. Here, we anticipate the statement of the main result in order to complete this general introduction of HyTNs.

Theorem 4 ([12])**.** *A HyTN* $\mathcal{H}$ *is consistent if and only if it is conservative. Moreover, when all weights are integral, then* $\mathcal{H}$ *admits an integral scheduling if and only if it is conservative.*

Let us now focus on multi-head hypergraphs and multi-head HyTNs.

Definition 6 (Hypergraph)**.** *A* hypergraph $\mathcal{H}$ *is a pair* $\langle V, \mathcal{A} \rangle$*, where* V *is the set of nodes, and* $\mathcal{A}$ *is the set of* hyperarcs*. Each hyperarc* $A = \langle t_A, H_A, w_A \rangle \in \mathcal{A}$ *has a distinguished node* t_A*, called the* tail *of* A*, and a nonempty set* $H_A \subseteq V \setminus \{t_A\}$ *containing the* heads *of* A*; to each head* $v \in H_A$ *is associated a* weight $w_A(v) \in \mathbf{Z}$*.*

Provided that $|A| \triangleq |H_A \cup \{t_A\}|$, the *size* of a hypergraph $\mathcal{H} = \langle V, \mathcal{A} \rangle$ is defined as $m_{\mathcal{A}} \triangleq \sum_{A \in \mathcal{A}} |A|$, and it is used as a measure for the encoding length of $\mathcal{H}$. If $|A| = 2$, then $A = \langle u, v, w \rangle$ can be regarded as a *standard arc*. In this way, hypergraphs generalize graphs.

At this point, a (multi-head) HyTN is thus a weighted hypergraph $\mathcal{H} = \langle V, \mathcal{A} \rangle$ where a node represents an *event* to be scheduled in time, and a hyperarc represents a set of temporal distance *constraints* between the *tail* and the *heads*.

In the (multi-head) HyTN framework the consistency problem is defined to be the following decision problem.

Definition 7 (HyTN-Consistency)**.** *Given a (multi-head) HyTN* $\mathcal{H} = \langle V, \mathcal{A} \rangle$*, decide whether there exists a scheduling function* $\phi : V \to \mathbf{R}$ *such that:*

$$\phi(t_A) \geq \min_{v \in H_A} \phi(v) - w_A(v), \ \forall \ A \in \mathcal{A}$$

any such scheduling $\phi : V \to \mathbf{R}$ *is called feasible.*

A HyTN is called consistent whenever it admits at least one feasible scheduling. The problem of checking whether a given HyTN is consistent is named HyTN-Consistency.

The computational equivalence between checking the consistency of (multi-head) HyTNs and determining the winning sets of MPGs was pointed out in [11, 12]. The tightest currently known worst-case time complexity for solving HyTN-Consistency is expressed by the following theorem, which was proved in [12] by relying on the Value-Iteration Algorithm for MPGs [5]. The approach was shown to be robust thanks to extensive experimental evaluations in [12], also see [4].

Theorem 5. *[11, 12] The following propositions hold on (multi-head) HyTNs.*

1. *There exists an $O((|V|+|\mathcal{A}|)m_{\mathcal{A}}W)$ pseudo-polynomial time algorithm for checking HyTN-Consistency;*
2. *There exists an $O((|V|+|\mathcal{A}|)m_{\mathcal{A}}W)$ pseudo-polynomial time algorithm such that, given in input any consistent HyTN $\mathcal{H}=(V,\mathcal{A})$, it returns to output a feasible scheduling $\phi : V \rightarrow \mathbf{R}$ of $\mathcal{H}$;*
3. *There exists an $O((|V|+|\mathcal{A}|)m_{\mathcal{A}}W)$ pseudo-polynomial time algorithm such that, given in input any inconsistent HyTN $\mathcal{H}=(V,\mathcal{A})$, it returns to output a (generalised) negative cycle $(S,\mathcal{C})$ of $\mathcal{H}$;*

Here, $W \triangleq \max_{A\in\mathcal{A},v\in H_A} |w_A(v)|$.

In the forthcoming section we shall turn our attention to *conditional* temporal planning; where we generalise Conditional Simple Temporal Networks (CSTNs) by introducing Conditional Hyper Temporal Networks (CHyTNs).

3 Conditional Simple/Hyper Temporal Networks

In order to provide a formal support to the present work, this section recalls the basic formalism, terminology and known results on CSTPs and CSTNs. Since the forthcoming definitions concerning CSTNs are mostly inherited from the literature, the reader is referred to [28] and [21] for an intuitive semantic discussion and for some clarifying examples of the very same CSTN model.

In fact, in 2003, [28] introduced the *Conditional Simple Temporal Problem (CSTP)* as an extension of standard temporal constraint-satisfaction models used in non-conditional temporal planning. A CSTP augments an STN to include *observation* events. Each observation event has a *boolean variable* (or *proposition*) associated with it. When the observation event is executed, the truth-value of its associated proposition becomes known. In addition, each event and each constraint has a *label* that restricts the scenarios in which it plays a role. Although not included in the formal definition, [28] discussed some supplementary *reasonability assumptions* that any well-defined CSTP must satisfy. Subsequently, those conditions have been analyzed and formalized in [21], leading to the sound notion of *Conditional Simple Temporal Network* (CSTN), which is now recalled.

Let P be a set of boolean variables, a *label* is any (possibly empty) conjunction of variables, or negations of variables, drawn from P. The *empty label* is denoted by λ. The *label universe* of P, denoted P^*, is the set of all (possibly empty) labels whose (positive or negative) literals are drawn from P. Two labels, ℓ_1 and ℓ_2, are called *consistent*, denoted[1] by *Con*(ℓ_1,ℓ_2), when $\ell_1 \wedge \ell_2$ is satisfiable. A label ℓ_1 *subsumes* a label ℓ_2, denoted[1] by *Sub*(ℓ_1,ℓ_2), when the implication $\ell_1 \Rightarrow \ell_2$ holds. We are now in the position to recall the definition of CSTNs.

Definition 8 (CSTNs)**.** *A* Conditional Simple Temporal Network (CSTN) *is a tuple $\langle V, A, L, \mathcal{O}, \mathcal{O}V, P\rangle$ where:*

- *V is a finite set of* events*; $P=\{p_1,\ldots,p_q\}$ (some $q \in \mathbf{N}$) is a finite set of* boolean variables *(or* propositions*);*

[1]The notation *Con*$(\cdot,\cdot)$ and *Sub*$(\cdot,\cdot)$ is inherited from [28, 21].

- A *is a set of* labeled temporal constraints (LTCs) *each having the form* $\langle v - u \leq w(u,v), \ell \rangle$*, where* $u, v \in V$*,* $w(u,v) \in \mathbf{Z}$*, and* $\ell \in P^*$*;*

- $L : V \to P^*$ *is a function that assigns a label to each event in* V*;* $\mathcal{O}V \subseteq V$ *is a finite set of* observation events*;* $\mathcal{O} : P \to \mathcal{O}V$ *is a bijection that associates a unique observation event* $\mathcal{O}(p) = \mathcal{O}_p$ *to each proposition* $p \in P$*;*

- *The following* reasonability assumptions *must hold:*

  *(*WD1*) for any labeled constraint* $\langle v - u \leq w, \ell \rangle \in A$ *the label* ℓ *is satisfiable and subsumes both* $L(u)$ *and* $L(v)$*; intuitively, whenever a constraint* $v - u \leq w$ *is required to be satisfied, then its endpoints* u *and* v *must be scheduled (sooner or later) by the planning agent;*

  *(*WD2*) for each* $p \in P$ *and each* $u \in V$ *such that either* p *or* $\neg p$ *appears in* $L(u)$*, we require:* $\text{Sub}(L(u), L(\mathcal{O}_p))$*, and* $\langle \mathcal{O}_p - u \leq -\epsilon, L(u) \rangle \in A$ *for some* $\epsilon > 0$*; intuitively, whenever a label* $L(u)$ *of an event node* u *contains proposition* p*, and* u *gets eventually scheduled, then the observation event* $\mathcal{O}_p$ *must be scheduled strictly before* u *by the planning agent.*

  *(*WD3*) for each labeled constraint* $\langle v - u \leq w, \ell \rangle$ *and* $p \in P$*, for which either* p *or* $\neg p$ *appears in* ℓ*, it holds that* $\text{Sub}(\ell, L(\mathcal{O}_p))$*; intuitively, assuming a required constraint contains proposition* p*, then the observation event* $\mathcal{O}_p$ *must be scheduled (sooner or later) by the planner.*

We are now in the position to introduce the *Conditional Hyper Temporal Network (CHyTN)*, which is a natural extension and generalization of both the CSTN and the HyTN model, obtained by "melting" them together.

Definition 9 (CHyTNs)**.** *A* Conditional Hyper Temporal Network (CHyTN) *is a tuple* $\langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P \rangle$*, where* $V, P, L, \mathcal{O}$ *and* $\mathcal{O}V$ *are defined as in CSTNs (see Definition 8), and where* $\mathcal{A}$ *is a set of* labeled temporal hyper constraints *each having one of the following forms:*

- $A = (t, h, w, \ell)$*, where* (t, h, w) *is a standard arc and* $\ell \in P^*$ *is a label. In this case,* A *is said to be a* standard *labeled temporal constraint.*

- $A = (t_A, H_A, w_A, L_{H_A})$*, where* (t_A, H_A, w_A) *is a* multi-head *hyperarc and* $L_{H_A} : H_A \to P^* : h \mapsto \ell_h$ *is is a function that assigns to each head of* A *a label in* P^**. In this case,* A *is said to be a* multi-head *labeled temporal constraint.*

- $A = (T_A, h_A, w_A, L_{T_A})$*, where* $A = (T_A, h_A, w_A)$ *is a* multi-tail *hyperarc and* $L_{T_A} : T_A \to P^* : t \mapsto \ell_t$ *is is a function that assigns to each tail of* A *a label in* P^**. In this case,* A *is said to be a* multi-tail *labeled temporal constraint.*

- *The following* reasonability assumptions *must hold:*

  *(*WD1'*) for any labeled constraint* A*:*

  - *if* $A = (t, h, w, \ell)$ *is a standard labeled temporal constraint, the label* ℓ *is satisfiable and subsumes both* $L(t)$ *and* $L(h)$*;*
  - *if* $A = (t_A, H_A, w_A, L_{H_A})$ *is a multi-head labeled temporal constraint, then for each* $h \in H_A$ *the label* $L_{H_A}(h)$ *is satisfiable and subsumes both* $L(t_A)$ *and* $L(h)$*;*
  - *if* $A = (T_A, h_A, w_A, L_{T_A})$ *is a multi-tail labeled temporal constraint, then for each* $t \in T_A$ *the label* $L_{T_A}(h)$ *is satisfiable and subsumes both* $L(t)$ *and* $L(h_A)$*;*

*(*WD2*) for each $p \in P$ and each $u \in V$ such that either p or $\neg p$ appears in $L(u)$, we require:* $\text{Sub}(L(u), L(\mathcal{O}_p))$*, and $\langle \mathcal{O}_p - u \leq -\epsilon, L(u)\rangle \in \mathcal{A}$ for some (small) $\epsilon > 0$; This is the same reasonability assumptions WD2 as defined for CSTNs.*

*(*WD3'*) for each labeled constraint $A \in \mathcal{A}$ and $p \in P$:*

- *if $A = (t, h, w, \ell)$ is a standard labeled temporal constraint and p or $\neg p$ appears in ℓ, then* $\text{Sub}(\ell, L(\mathcal{O}_p))$*;*
- *if $A = (t_A, H_A, w_A, L_{H_A})$ is a multi-head labeled temporal constraint and either p or $\neg p$ appears in $L_{H_A}(h)$ for some $h \in H_A$, then* $\text{Sub}(L_{H_A}(h), L(\mathcal{O}_p))$*;*
- *if $A = (T_A, h_A, w_A, L_{T_A})$ is a multi-tail labeled temporal constraint and either p or $\neg p$ appears in $L_{T_A}(t)$ for some $t \in T_A$, then* $\text{Sub}(L_{T_A}(t), L(\mathcal{O}_p))$*;*

Of course, every CSTN is a CHyTN (i.e., one having standard labeled temporal constraints only).

Example 1. *Fig. 2a depicts an example of a CSTN $\Gamma_1 = \langle V, A, L, \mathcal{O}, \mathcal{O}V, P\rangle$ having three event nodes A, B and C as well as two observation events $\mathcal{O}_p$ and $\mathcal{O}_q$. Formally, $V = \{A, B, C, \mathcal{O}_p, \mathcal{O}_q\}$, $P = \{p, q\}$, $\mathcal{O}V = \{\mathcal{O}_p, \mathcal{O}_q\}$, $L(v) = \lambda$ for every $v \in V$, $\mathcal{O}(p) = \mathcal{O}_p, \mathcal{O}(q) = \mathcal{O}_q$. Moreover, the set of labeled temporal constraints is: $A = \{\langle C - A \leq 10, \lambda\rangle, \langle A - C \leq -10, \lambda\rangle, \langle B - A \leq 3, p \wedge \neg q\rangle, \langle A - B \leq 0, \lambda\rangle, \langle \mathcal{O}_p - A \leq 5, \lambda\rangle, \langle A - \mathcal{O}_p \leq 0, \lambda\rangle, \langle \mathcal{O}_q - A \leq 9, \lambda\rangle, \langle A - \mathcal{O}_q \leq 0, \lambda\rangle, \langle C - B \leq 2, q\rangle, \langle C - \mathcal{O}_p \leq 10, \lambda\rangle, \langle C - \mathcal{O}_q \leq 1, \neg p\rangle\}$.*

Fig. 2b depicts an example of a CHyTN $\Gamma_2 = \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P\rangle$. Notice that $V, L, \mathcal{O}, \mathcal{O}V$ and P are the same as in the CSTN Γ_1, whereas $\mathcal{A}$ is defined as follows: $\mathcal{A} = A \cup \{\alpha \triangleq (B, \{C, \mathcal{O}_q\}, \langle w_\alpha(C), w_\alpha(\mathcal{O}_q)\rangle = \langle 2, -1\rangle, \langle L_\alpha(C), L_\alpha(\mathcal{O}_q)\rangle = \langle \lambda, \lambda\rangle)\}$, where A is the set of (standard) labeled temporal constraints of the CSTN Γ_1 and the additional constraint α is a multi-head labeled temporal constraint with tail B and heads $C, \mathcal{O}_q$.

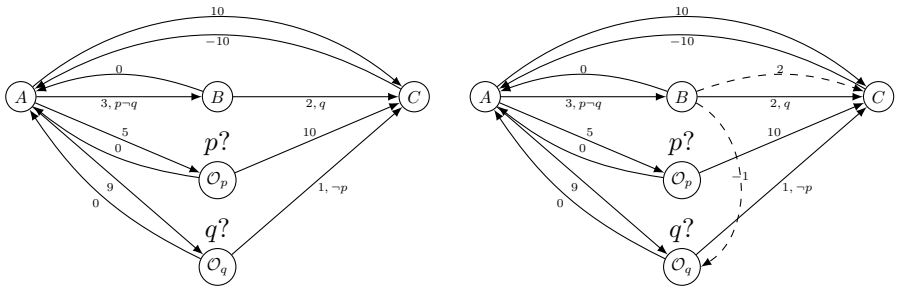

(a) A CSTN example Γ_1. (b) A CHyTN example Γ_2

Figure 2: An example of CSTN (a) and an example CHyTN (b).

In the following definitions we will implicitly refer to some CHyTN which is denoted $\Gamma = \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P\rangle$.

Definition 10 (Scenario)**.** *A* scenario *over a set P of boolean variables is a truth assignment $s : P \to \{\top, \bot\}$, i.e., s is a function that assigns a truth value to each proposition $p \in P$. The set of all scenarios over P is denoted Σ_P. If $s \in \Sigma_P$ is a scenario and $\ell \in P^*$ is a label, then $s(\ell) \in \{\top, \bot\}$ denotes the truth value of ℓ induced by s in the natural way.*

Notice that any scenario $s \in \Sigma_P$ can be described by means of the label $\ell_s \triangleq l_1 \wedge \cdots \wedge l_{|P|}$ such that, for every $1 \leq i \leq |P|$, the literal $l_i \in \{p_i, \neg p_i\}$ satisfies $s(l_i) = \top$.

Example 2. *Consider the set of propositional variables $P = \{p, q\}$. The scenario $s : P \to \{\top, \bot\}$ defined as $s(p) = \top$ and $s(q) = \bot$ can be compactly described by the label $\ell_s = p \wedge \neg q$.*

Definition 11 (Scheduling)**.** *A* scheduling *for a subset of events $U \subseteq V$ is a function $\phi : U \to \mathbf{R}$ that assigns a real number to each event in U. The set of all schedules over U is denoted Φ_U.*

Definition 12 (Scenario Restriction)**.** *Let $s \in \Sigma_P$ be a scenario. The* restriction *of V and $\mathcal{A}$ w.r.t. s are defined as follows:*

$$V_s^+ \triangleq \{v \in V \mid s(L(v)) = \top\};$$

$$\begin{aligned}\mathcal{A}_s^+ \triangleq \Big\{(u,v,w) \mid \exists(\ell \in P^*) \text{ s.t. } (u,v,w,\ell) \in \mathcal{A} \text{ and } s(\ell) = \top\Big\} \cup \\ \cup\Big\{(t, H'_A, w'_A) \mid \exists(H_A \supseteq H'_A; L_{H_A} : H_A \to P^*; w_A : H_A \to \mathbf{Z}) \text{ s.t. } (t, H_A, w_A, L_{H_A}) \in \mathcal{A}, \\ w'_A = w_{A|_{H'_A}}, \forall(h \in H_A)\, s(L_{H_A}(h)) = \top \iff h \in H'_A\Big\} \cup \\ \cup\Big\{(T'_A, h, w'_A) \mid \exists(T_A \supseteq T'_A; L_{T_A} : T_A \to P^*; w_A : T_A \to \mathbf{Z}) \text{ s.t. } (T_A, h, w_A, L_{T_A}) \in \mathcal{A}, \\ w'_A = w_{A|_{T'_A}}, \forall(t \in T_A)\, s(L_{T_A}(t)) = \top \iff t \in T'_A\Big\}.\end{aligned}$$

The restriction of Γ w.r.t. s is defined as $\Gamma_s^+ \triangleq \langle V_s^+, \mathcal{A}_s^+\rangle$.
Finally, it is worth to introduce the notation $V_{s_1,s_2}^+ \triangleq V_{s_1}^+ \cap V_{s_2}^+$.

Notice that the restriction Γ_s^+ of a CHyTN Γ is always a HyTN, whereas Γ_s^+ is an STN whenever Γ is a CSTN.

Example 1. *Fig. 3 depicts the restriction STN ${\Gamma_1}_s^+$ of the CSTN Γ_1, and the restriction HyTN ${\Gamma_2}_s^+$ of the CHyTN Γ_2 (see Example 1 and Fig. 2), with respect to the scenario defined as: $s(p) = \bot$ and $s(q) = \top$.*

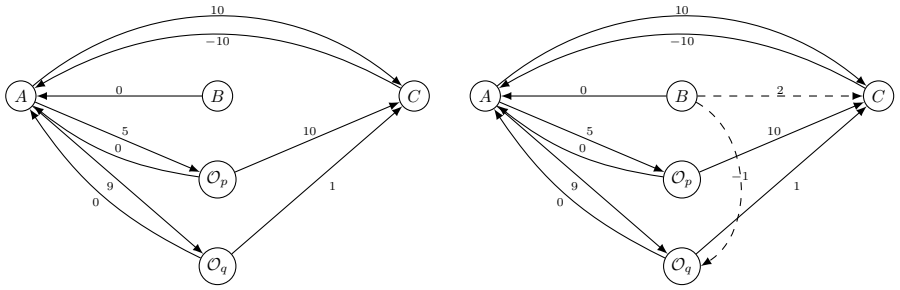

(a) The restriction STN ${\Gamma_1}_s^+$ of the CSTN Γ_1 w.r.t. $s(p) = \bot$ and $s(q) = \bot$

(b) The restriction HyTN ${\Gamma_2}_s^+$ of the CHyTN Γ_2 w.r.t. $s(p) = \bot$ and $s(q) = \bot$

Figure 3: The restriction ${\Gamma_1}_s^+$ of Γ_1 (a), and the restriction ${\Gamma_2}_s^+$ of Γ_2 (b), with respect to the scenario $s(p) = \bot$ and $s(q) = \top$

Definition 13 (Execution Strategy)**.** *An* execution strategy *for Γ is a mapping $\sigma : \Sigma_P \to \Phi_V$ such that, for any scenario $s \in \Sigma_P$, the domain of the scheduling $\sigma(s)$ is V_s^+. The set of execution strategies of Γ is denoted by $\mathcal{S}_\Gamma$. The* execution time *of an event node $v \in V_s^+$ in the schedule $\sigma(s) \in \Phi_{V_s^+}$ is denoted by $[\sigma(s)]_v$.*

Definition 14 (Scenario History)**.** *Let $\sigma \in \mathcal{S}_\Gamma$ be an execution strategy, let $s \in \Sigma_P$ be a scenario and let $v \in V_s^+$ be an event node. The* scenario history *scHist(v, s, σ) of v in the scenario s for the strategy σ is defined as follows:*

$$\textit{scHist}(v, s, \sigma) \triangleq \Big\{\big(p, s(p)\big) \mid p \in P,\, \mathcal{O}_p \in V_s^+ \cap \mathcal{O}V,\, [\sigma(s)]_{\mathcal{O}_p} < [\sigma(s)]_v\Big\}.$$

The scenario history can be compactly expressed by the conjunction of the literals corresponding to the observations comprising it. Thus, we may treat a scenario history as though it were a label.

Definition 15 (Viable Execution Strategy)**.** *We say that $\sigma \in \mathcal{S}_\Gamma$ is a* viable *execution strategy if, for each scenario $s \in \Sigma_P$, the scheduling $\sigma(s) \in \Phi_V$ is feasible for the restriction HyTN (or STN) Γ_s^+.*

Definition 16 (Dynamic Consistency)**.** *An execution strategy $\sigma \in \mathcal{S}_\Gamma$ is called* dynamic *if, for any $s_1, s_2 \in \Sigma_P$ and any event node $v \in V^+_{s_1,s_2}$, the following implication holds:*

$$\text{Con}(\text{scHst}(v, s_1, \sigma), s_2) \Rightarrow [\sigma(s_1)]_v = [\sigma(s_2)]_v. \tag{DC}$$

We say that Γ is dynamically-consistent *if it admits $\sigma \in \mathcal{S}_\Gamma$ which is both viable and dynamic.*

Definition 17 (DC-Checking)**.** *The problem of checking whether a given CHyTN (which allows* both *multi-head and multi-tail labeled temporal constraints) is dynamically-consistent is named* General-CHyTN-DC.

The problem of checking whether a given CHyTN allowing only *multi-head* or only *multi-tail labeled temporal constraints, is dynamically-consistent is named* CHyTN-DC.

The problem of checking whether a given CSTN is dynamically-consistent is named CSTN-DC.

Example 3. *Consider the CHyTN Γ of Fig. 2b, and let the scenarios s_1, s_2, s_3, s_4 be defined as: $s_1(p) = \bot$, $s_1(q) = \bot$; $s_2(p) = \bot$, $s_2(q) = \top$; $s_3(p) = \top$, $s_3(q) = \bot$; $s_4(p) = \top$, $s_4(q) = \top$. It follows an example of execution strategy $\sigma \in \mathcal{S}_\Gamma$: $[\sigma(s_i)]_A = 0$ for every $i \in \{1,2,3,4\}$; $[\sigma(s_i)]_B = 8$ for every $i \in \{1,2,4\}$ and $[\sigma(s_3)]_B = 3$; $[\sigma(s_i)]_C = 10$ for every $i \in \{1,2,3,4\}$; $[\sigma(s_i)]_{\mathcal{O}_p} = 1$ for every $i \in \{1,2,3,4\}$. $[\sigma(s_i)]_{\mathcal{O}_q} = 2$ for every $i \in \{3,4\}$ and $[\sigma(s_i)]_{\mathcal{O}_q} = 9$ for every $i \in \{1,2\}$. The reader can check that σ is viable and dynamic. Indeed, σ admits the tree-like representation depicted in Fig 4.*

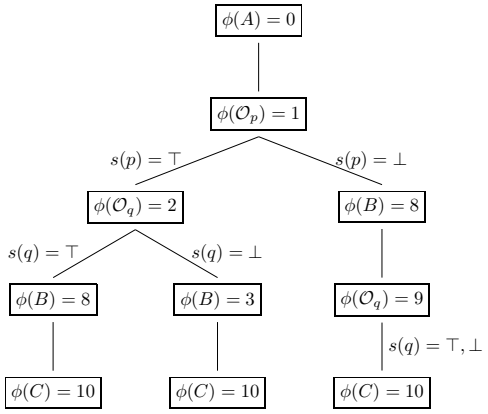

Figure 4: A tree-like representation of a dynamic execution strategy σ for the CHyTN Γ of Fig. 2b, where s denotes scenarios and $\phi(X) = [\sigma(s)]_X$ denotes the corresponding scheduling.

We now introduce next a crucial notion for studying dynamic-consistency of CSTNs, the *difference set* $\Delta(s_1; s_2)$.

Definition 18 (Difference Set)**.** *Let $s_1, s_2 \in \Sigma_P$ be two scenarios. The set of observation events in $V^+_{s_1} \cap \mathcal{O}V$ at which s_1 and s_2 differ is denoted by $\Delta(s_1; s_2)$. Formally,*

$$\Delta(s_1; s_2) \triangleq \{\mathcal{O}_p \in V^+_{s_1} \cap \mathcal{O}V \mid s_1(p) \neq s_2(p)\}.$$

Notice that commutativity may not hold, i.e., in general it may be the case that $\Delta(s_1; s_2) \neq \Delta(s_2; s_1)$.

Example 4. *Consider the CSTN Γ of Fig. 2a and the scenarios s_1, s_2 defined as follows: $s_1(p) = \bot$, $s_1(q) = \bot$; $s_2(p) = \bot$, $s_2(q) = \top$. Then, $\Delta(s_1; s_2) = \{\mathcal{O}_q\}$.*

The next lemma will be useful later on in Section 4.

Lemma 4. *Let $s_1, s_2 \in \Sigma_P$ and $v \in V^+_{s_1,s_2}$. Let $\sigma \in \mathcal{S}_\Gamma$ be an execution strategy. Then, σ is dynamic if and only if the following implication holds for every $s_1, s_2 \in \Sigma_P$ and for every $u \in V^+_{s_1,s_2}$:*

$$\Big(\bigwedge_{v \in \Delta(s_1;s_2)} [\sigma(s_1)]_u \leq [\sigma(s_1)]_v \Big) \Rightarrow [\sigma(s_1)]_u = [\sigma(s_2)]_u \tag{L4}$$

Proof. Notice that, by definition of *Con*$(\cdot,\cdot)$ and *scHst*$(\cdot,\cdot,\cdot)$, *Con*(*scHst*$(u, s_1, \sigma), s_2)$ holds if and only if there is no observation event $v \in \Delta(s_1; s_2)$ which is scheduled by $\sigma(s_1)$ strictly before u. Therefore, *Con*(*scHst*$(u, s_1, \sigma), s_2)$ holds if and only if $\bigwedge_{v \in \Delta(s_1;s_2)} [\sigma(s_1)]_u \leq [\sigma(s_1)]_v$. At this point, substituting the *Con*(*scHst*$(u, s_1, \sigma), s_2)$ expression with the equivalent formula $\bigwedge_{v \in \Delta(s_1;s_2)} [\sigma(s_1)]_u \leq [\sigma(s_1)]_v$ inside the definition of dynamic execution strategy (see Definition 16), then the thesis follows. □

4 Algorithmics of Dynamic-Consistency

To start with, we offer the following coNP-hardness lower bound on CSTN-DC.

Since any CSTN is also a CHyTN, the same coNP-hardness result holds on CHyTN-DC and General-CHyTN-DC.

Theorem 6. *CSTN-DC is a* coNP*-hard problem even if input instances $\Gamma = \langle V, A, L, \mathcal{O}, \mathcal{OV}, P\rangle$ are restricted to satisfy $w_A(\cdot) \in \{-1, 0\}$ and $\ell \in \{p, \neg p \mid p \in P\} \cup \{\lambda\}$ for every $(u, v, w, \ell) \in A$.*

Proof. We reduce 3-SAT to the complement of CSTN-DC. Let φ be a boolean formula in 3CNF. Let X be the set of variables and let $\mathcal{C} = \{C_0, \ldots, C_{m-1}\}$ be the set of clauses comprising $\varphi = \bigwedge_{j=0}^{m-1} C_j$.

(1) Let N^φ be the CSTN $\langle V^\varphi, A^\varphi, L^\varphi, \mathcal{O}^\varphi, \mathcal{OV}^\varphi, P^\varphi\rangle$, where: $V^\varphi \triangleq X \cup \mathcal{C}$, and all the nodes are given empty label, i.e., $L^\varphi(v) = \lambda$ for every $v \in V^\varphi$; $P^\varphi \triangleq \mathcal{OV}^\varphi \triangleq X$; $\mathcal{O}^\varphi$ is the identity function; for every $u, v \in \mathcal{OV}^\varphi$ we have $\langle u - v \leq 0, \lambda\rangle \in A^\varphi$; for every $x \in X$ and $C \in \mathcal{C}$ we have $\langle x - C \leq -1, \lambda\rangle \in A^\varphi$; for each $j = 0, \ldots, m-1$ and for each literal $\ell \in C_j$, we have $\langle C_j - C_{(j+1) \bmod m} \leq -1, \ell\rangle \in \mathcal{A}_\varphi$. Notice that $|V^\varphi| = n + m$ and $|A^\varphi| = n^2 + nm + 3m$.

(2) Assume that φ is satisfiable. Let ν be a satisfying truth-assignment of φ. In order to prove that N^φ is not dynamically-consistent, observe that the restriction of N^φ w.r.t. the scenario ν is a non-consistent STN. Indeed, if for every $j = 0, \ldots, m-1$ we pick a standard arc $\langle C_j - C_{(j+1) \bmod m} \leq -1, \ell_j\rangle$ with ℓ_j being a literal in C_j such that $\nu(\ell_j) = \top$, then we obtain a negative circuit.

(3) Assume that φ is unsatisfiable. In order to prove that N^φ is dynamically-consistent, we exhibit a viable and dynamic execution strategy σ for N^φ. First, schedule every $x \in X$ at $\sigma(x) \triangleq 0$. Therefore, by time 1, the planner has full knowledge of the observed scenario ν. Since φ is unsatisfiable, there exists an index j_ν such that $\nu(C_{j_\nu}) = \bot$. At this point, set $\sigma(C_{(j_\nu+k) \bmod m}) \triangleq k$ for $k = 1, \ldots, m$. The reader can verify that σ is viable and dynamic for N^φ. □

An illustration of the CSTN N^φ, which was constructed in the proof of Theorem 6, is shown in Fig. 5.

It remains currently open whether or not CSTN-DC lies in PSPACE and whether or not it is PSPACE-hard. However, now we can show that when the input CHyTN instances are allowed to have *both* multi-heads *and* multi-tail LTCs then the corresponding DC-Checking problem becomes PSPACE-hard.

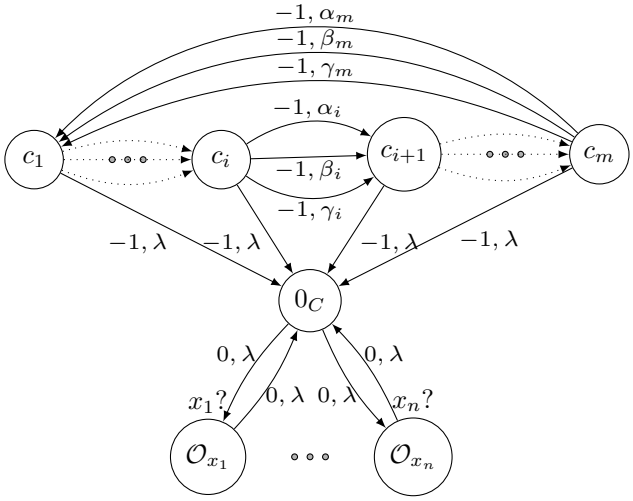

Figure 5: The CSTN $C(\varphi)$ where $\varphi(x_1, \ldots, x_n) = \bigwedge_{i=1}^{m} c_i$ for $c_i = (\alpha_i \vee \beta_i \vee \gamma_i)$.

Theorem 7. *General-CHyTN-DC is a* PSPACE*-hard problem, even if the input instances* $\Gamma = \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P\rangle$ *are restricted to satisfy the following two constraints:*

– $w_a(\cdot) \in [-n-1, n+1] \cap \mathbf{Z}$ *and* $\ell_a \in \{p, \neg p \mid p \in P\} \cup \{\lambda\}$ *for every weight* w_a *and label* ℓ_a *appearing in any* standard *LTC* $a \in \mathcal{A}$*;*

– $w_A(\cdot) \in \{-1, 0, 1\}$, $\ell_A = \lambda$ *and* $|A| \leq 2$ *for every weight* w_A *and label* ℓ_A *appearing in any* multi-tail/head *LTC* $A \in \mathcal{A}$*.*

Proof. To show that General-CHyTN-DC is PSPACE-hard, we describe a reduction from the problem 3-CNF-TQBF (True Quantified Boolean Formula in 3-CNF).

Let us consider a 3-CFN quantified Boolean formula with $n \geq 1$ variables and $m \geq 1$ clauses:

$$\varphi(x_1, \ldots, x_n) = Q_1 x_1 \ldots Q_n x_n \bigwedge_{i=1}^{m} (\alpha_i \vee \beta_i \vee \gamma_i),$$

where for every $j \in [n]$ the symbol Q_j is either $\exists$ or $\forall$, and where $\mathcal{C}_i = (\alpha_i \vee \beta_i \vee \gamma_i)$ is the i-th clause of φ and each $\alpha_i, \beta_i, \gamma_i \in \{x_j, \overline{x}_j \mid 1 \leq j \leq n\}$ is a positive or negative literal. We also say that $Q_1 x_1 \ldots Q_n x_n$ is the *prefix* of φ.

We associate to φ a CHyTN $\Gamma_\varphi = \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P\rangle$. In so doing, our first goal is that to simulate the interaction between two players, i.e., Player-$\exists$ (that corresponds to the Planner, in CHyTNs) and Player-$\forall$ (that corresponds to Nature, in CHyTNs), which corresponds directly to the chain of alternating quantifiers in the prefix of φ: naturally, the Planner is going to control those variables that are quantified existentially in φ, whereas Nature is going to control (by means of some observation events in $\mathcal{O}V$) those variables that are quantified universally in φ.

To begin, we let $P := \{x_j \mid$ "$\forall x_j$" appears in the prefix of $\varphi\}$. Moreover, V contains a node z (i.e., the *zero* node) that represents the reference initial node for Γ_φ, i.e., node z has to be executed at time 0.

Next, for each existential quantification $\exists x_j$ in the prefix of φ, V contains a node named t_{x_j} and $\mathcal{A}$ contains the following two standard LTCs: $(z, t_{x_j}, j+1, \lambda)$ and $(t_{x_j}, z, -j, \lambda)$; the underlying intuition being that, during the execution, it will be the responsibility of the planner to schedule t_{x_j} either at time j (and this will mean that the Planner chooses to set x_j to `false`) or at time $j+1$ (and this will mean that the Planner chooses to set x_j to `true`). See Fig. 6a for an illustration of the $\exists x_j$'s gadget.

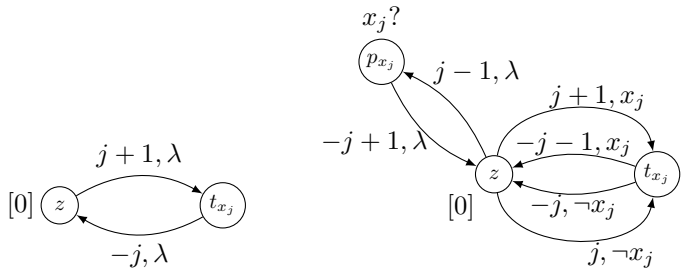

(a) Gadget for a 3-CNF-TQBF existentially quantified variable $\exists x_j$.

(b) Gadget for a 3-CNF-TQBF universally quantified variable $\forall x_j$.

Figure 6: Gadgets for quantified variables used in the reduction from 3-CNF-TQBF to General-CHyTN-DC.

Moreover, for each universal quantification $\forall x_j$ in the prefix of φ (i.e., for each $x_j \in P$), V contains two nodes named p_{x_j} and t_{x_j}. In particular, p_{x_j} is an observation node (i.e., $p_{x_j} \in \mathcal{O}V$) such that $\mathcal{O}(x_j) = p_{x_j}$; actually $\mathcal{O}V := \{p_{x_j} \mid x_j \in P\}$. Also, for each such $\forall x_j$ in φ's prefix (i.e., for each $x_j \in P$), $\mathcal{A}$ contains the following six standard LTCs: $(z, p_{x_j}, j-1, \lambda)$, $(p_{x_j}, z, -j+1, \lambda)$, and then $(z, t_{x_j}, j+1, x_j)$, $(t_{x_j}, z, -j-1, x_j)$, also $(z, t_{x_j}, j, \neg x_j)$, $(t_{x_j}, z, -j, \neg x_j)$; the underlying intuition being that it will be this time the Nature the one to choose whether to schedule t_{x_j} at time j (i.e., setting x_j to `false` by controlling the observation event p_{x_j}) or at time $j+1$ (i.e., setting x_j to `true`, again, by controlling the observation event p_{x_j}). Fig. 6b provides an illustration of the $\forall x_j$'s gadget.

In both cases, i.e., for existential $\exists x_j$ and universal $\forall x_j$ quantified variables, notice that the weights of the involved standard temporal constraints depend on j in such a way that their scheduling times and their corresponding propositional choices really have the possibility to occur one after the other in time; more precisely, for every $j \in [n]$, t_{x_j} is going to be scheduled either at time j (and this corresponds to the case in which x_j is `true`) or at time $j+1$ (when instead x_j is `false`); in addition to this, when x_j is quantified universally in φ (i.e., when $x_j \in P$), then the observation event that determines its propositional value (i.e., p_{x_j}) is always scheduled at time $j-1$ (and this leaves space widely for the reaction time; actually, an entire unit of time between time $j-1$ and time j).

This concludes the description of our gadgets for simulating the chain of alternating quantifiers in the prefix of φ.

At this point, we shall describe two additional gadgets for simulating the 3-CNF formula $\bigwedge_{i=1}^{m}(\alpha_i \vee \beta_i \vee \gamma_i)$ of φ, one gadget is for the literals and the second one is for the clauses. We have an additional node in V, named z', which is always scheduled at time $n+1$; for this, $\mathcal{A}$ contains the following two standard LTCs: $(z, z', n+1, \lambda)$ and $(z', z, -n-1, \lambda)$. Next, we have a gadget for the positive (i.e., x_j) and the negative (i.e., $\neg x_j$) literals. It goes as follows: for each $j \in [n]$, V contains two nodes named l_{x_j} (i.e., positive literal) and $l_{\overline{x}_j}$ (i.e., negative literal); moreover, $\mathcal{A}$ contains the following four standard LTCs, $(z', l_{x_j}, 1, \lambda)$, $(l_{x_j}, z', 0, \lambda)$ and $(z', l_{\overline{x}_j}, 1, \lambda)$, $(l_{\overline{x}_j}, z', 0, \lambda)$, plus the following *multi-head* LTC,

$$A^h(l_{x_j}, l_{\overline{x}_j}) := \Big(z', \{l_{x_j}, l_{\overline{x}_j}\}, \langle w(l_{x_j}), w(l_{\overline{x}_j})\rangle = \langle 0, 0\rangle, \langle L(l_{x_j}), L(l_{\overline{x}_j})\rangle = \langle \lambda, \lambda\rangle\Big),$$

and the following *multi-tail* LTC,

$$A^t(l_{x_j}, l_{\overline{x}_j}) := \Big(\{l_{x_j}, l_{\overline{x}_j}\}, z', \langle w(l_{x_j}), w(l_{\overline{x}_j})\rangle = \langle -1, -1\rangle, \langle L(l_{x_j}), L(l_{\overline{x}_j})\rangle = \langle \lambda, \lambda\rangle\Big).$$

The idea here is that the standards LTCs are going to enforce the scheduling time of both l_{x_j} and $l_{\overline{x}_j}$ within the real interval $[n+1, n+2]$ (i.e., not before z' and at most 1 time unit after z'). Meanwhile, the multi-head

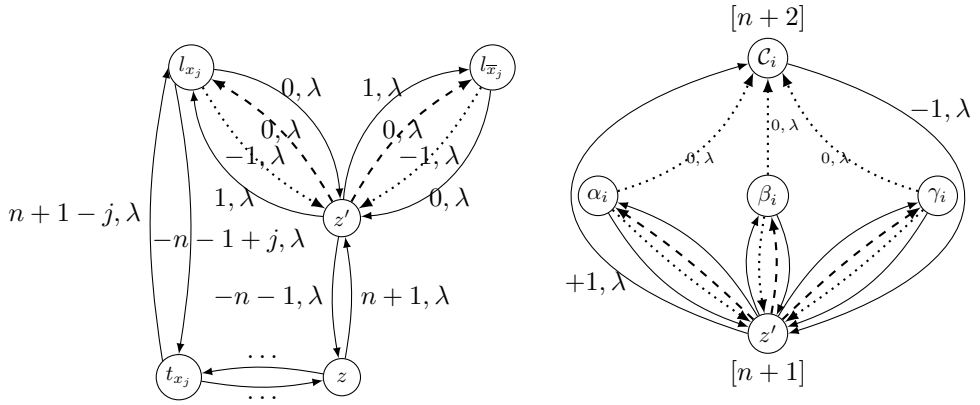

(a) Gadget for the 3-CNF-TQBF positive x_j and negative literal $\neg x_j$.

(b) Gadget for 3-CNF-TQBF clause $\mathcal{C}_i = (\alpha_i \vee \beta_i \vee \gamma_i)$ where each $\alpha_i, \beta_i, \gamma_i$ is a positive or negative literal.

Figure 7: Gadgets used for variables and clauses in the reduction from 3-CNF-TQBF to General-CHyTN-DC.

constraint $A^h(l_{x_j}, l_{\overline{x}_j})$ is going to enforce that at least one between l_{x_j} and $l_{\overline{x}_j}$ happen not later than time $n+1$ (i.e., not later than the scheduling time of z'); similarly, the multi-tail constraint $A^t(l_{x_j}, l_{\overline{x}_j})$ is going to enforce that at least one between l_{x_j} and $l_{\overline{x}_j}$ happen not before time $n+2$ (i.e., not before the scheduling time of z' plus 1). Therefore, exactly one between l_{x_j} and $l_{\overline{x}_j}$ will be enforced to happen at time $n+1$, and the other one at time $n+2$.

Got to this point, the key idea is that, for every $j \in [n]$, we can enforce the scheduling time of each node l_{x_j} and $l_{\overline{x}_j}$ to be uniquely determined, according to a suitable translation of the scheduling time of t_{x_j}. In particular, we want to schedule at time $n+1$ (i.e., at the same scheduling time of z') the one node between l_{x_j} and $l_{\overline{x}_j}$ whose corresponding literal was chosen to be `false` (that is l_{x_j} if t_{x_j} was scheduled at time j, and that is $l_{\overline{x}_j}$ if t_{x_j} was scheduled at time $j+1$); similarly, we want to schedule at time $n+2$ (i.e., at the same scheduling time of z' *plus* 1 time unit) the one node between l_{x_j} and $l_{\overline{x}_j}$ whose corresponding literal was chosen to be `true` (that is l_{x_j} if t_{x_j} was scheduled at time $j+1$, and that is $l_{\overline{x}_j}$ if t_{x_j} was scheduled at time j). In order to achieve this setting, for each $j \in [n]$, $\mathcal{A}$ contains the following two standard LTCs: $(t_{x_j}, l_{x_j}, n+1-j, \lambda)$ and $(l_{x_j}, t_{x_j}, -n-1+j, \lambda)$; in this way, l_{x_j} is enforced to happen at the same time of t_{x_j} *plus* $n+1-j$ units of time. Therefore, if t_{x_j} was scheduled at time j (i.e., x_j is `false`), then node l_{x_j} is scheduled at time $j+n+1-j = n+1$; otherwise, if t_{x_j} was scheduled at time $j+1$ (i.e., x_j is `true`), then node l_{x_j} is scheduled at time $j+1+n+1-j = n+2$. At this point, the scheduling time of the node $l_{\overline{x}_j}$ is determined uniquely thanks to the hyperarcs $A^h(l_{x_j}, l_{\overline{x}_j}), A^t(l_{x_j}, l_{\overline{x}_j})$ and the standard constraints $(z', l_{\overline{x}_j}, 1, \lambda)$, $(l_{\overline{x}_j}, z', 0, \lambda)$: If the node l_{x_j} is scheduled at time $n+1$ (i.e., if x_j is `false`), then $l_{\overline{x}_j}$ must be scheduled at time $n+1+1 = n+2$ (i.e., if $\overline{x}_j$ is `true`) so that to satisfy $A^t(l_{x_j}, l_{\overline{x}_j})$ and $(z', l_{\overline{x}_j}, 1, \lambda)$; otherwise, if l_{x_j} is scheduled at time $n+2$ (i.e., if x_j is `true`), then $l_{\overline{x}_j}$ must be scheduled at time $n+1+0 = n+1$ (i.e., if $\overline{x}_j$ is `false`) so that to satisfy $A^h(l_{x_j}, l_{\overline{x}_j})$ and $(l_{\overline{x}_j}, z', 0, \lambda)$.

Finally we have come to describe the clauses' gadget: For each $i \in [m]$, the CHyTN Γ_φ contains a node $\mathcal{C}_i$ for each clause $\mathcal{C}_i = (\alpha_i \vee \beta_i \vee \gamma_i)$ of φ; also, each node $\mathcal{C}_i$ is connected by:

– a multi-tail hyperarc with head in $\mathcal{C}_i$ and tails over the literals $\alpha_i, \beta_i, \gamma_i$ occurring in $\mathcal{C}_i$ and having weight 0 and label λ, i.e., by a multi-tail LTC:

$$A^c(\alpha_i, \beta_i, \gamma_i) := \Big(\{\alpha_i, \beta_i, \gamma_i\}, \mathcal{C}_i, \langle w(\alpha_i), w(\beta_i), w(\gamma)_i\rangle = \langle 0,0,0\rangle, \langle L(\alpha_i), L(\beta_i), L(\gamma)_i\rangle = \langle \lambda, \lambda, \lambda\rangle\Big),$$

for some literals $\alpha_i, \beta_i, \gamma_i \in \{l_{x_j}, l_{\overline{x}_j} \mid 1 \leq j \leq n\}$.

– two standard and opposite LTCs, $(z', C_i, 1, \lambda)$ and $(C_i, z', -1, \lambda)$, with node z'.

See Fig. 7b for an illustration of the clauses' gadget. Every node of Γ_φ has an empty label, i.e., $L(v) = \lambda$ for every $v \in V$. This concludes our prose description of Γ_φ.

More formally and succinctly, the CHyTN $\Gamma_\varphi = \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P\rangle$ is defined as follows:

- $P := \{x_j \mid$ "$\forall x_j$" appears in the prefix of $\varphi\}$;
- – $V := \{z, z'\} \cup \{t_{x_j} \mid 1 \leq j \leq n\} \cup \{p_{x_j} \mid x_j \in P\} \cup$
  $\cup \{l_{x_j} \mid 1 \leq j \leq n\} \cup \{l_{\overline{x}_j} \mid 1 \leq j \leq n\} \cup \{\mathcal{C}_i \mid 1 \leq i \leq m\}$;
  – $\mathcal{O}V := \{p_{x_j} \mid x_j \in P\}$ and $\mathcal{O}(x_j) = p_{x_j}$ for every $x_j \in P$;
  – $L(v) = \lambda$ for every $v \in V$;
- $\mathcal{A} := \bigcup_{j:\text{“}\exists x_j\text{”}\in\varphi} \exists\text{-Qnt}_j \cup \bigcup_{j:\text{“}\forall x_j\text{”}\in\varphi} \forall\text{-Qnt}_j \cup \bigcup_{j=1}^{n} \text{Var}_j \cup \bigcup_{i=1}^{m} \text{Cla}_i \cup \big\{(z, z', n+1, \lambda), (z', z, -n-1, \lambda)\big\}$,
  where:
  – $\exists\text{-Qnt}_j := \Big\{(z, t_{x_j}, j+1, \lambda), (t_{x_j}, z, -j, \lambda)\Big\}$;
  This defines the existential quantifier gadget as depicted in Fig. 6a;
  – $\forall\text{-Qnt}_j := \Big\{(z, p_{x_j}, j-1, \lambda), (p_{x_j}, z, -j+1, \lambda),$
  $(z, t_{x_j}, j+1, x_j), (t_{x_j}, z, -j-1, x_j), (z, t_{x_j}, j, \neg x_j), (t_{x_j}, z, -j, \neg x_j)\Big\}$;
  This defines the universal quantifier gadget as depicted in Fig. 6b;
  – $\text{Var}_j = \Big\{(z', l_{x_j}, 1, \lambda), (l_{x_j}, z', 0, \lambda), (z', l_{\overline{x}_j}, 1, \lambda), (l_{\overline{x}_j}, z', 0, \lambda),$
  $A_j^t := \Big(\{l_{x_j}, l_{\overline{x}_j}\}, z', \langle w_{A_j^t}(l_{x_j}), w_{A_j^t}(l_{\overline{x}_j})\rangle = \langle -1, -1\rangle, \langle L_{A_j^t}(l_{x_j}), L_{A_j^t}(l_{\overline{x}_j})\rangle = \langle \lambda, \lambda\rangle\Big),$
  $A_j^h := \Big(z', \{l_{x_j}, l_{\overline{x}_j}\}, \langle w_{A_j^h}(l_{x_j}), w_{A_j^h}(l_{\overline{x}_j})\rangle = \langle 0, 0\rangle, \langle L_{A_j^h}(l_{x_j}), L_{A_j^h}(l_{\overline{x}_j})\rangle = \langle \lambda, \lambda\rangle\Big),$
  $(t_{x_j}, l_{x_j}, n+1-j, \lambda), (l_{x_j}, t_{x_j}, -n-1+j, \lambda)\Big\}$.
  This defines the variable gadget for x_j as depicted in Fig. 7a;
  – $\text{Cla}_i = \Big\{(z', \mathcal{C}_j, 1), (\mathcal{C}_j, z', -1),$
  $A_i^c := \Big(\{\alpha_j, \beta_j, \gamma_j\}, \mathcal{C}_j, \langle w_{A_i^c}(\alpha_j), w_{A_i^c}(\beta_j), w_{A_i^c}(\gamma_j)\rangle = \langle 0, 0, 0\rangle,$
  $\langle L_{A_i^c}(\alpha_j), L_{A_i^c}(\beta_j), L_{A_i^c}(\gamma_j)\rangle = \langle \lambda, \lambda, \lambda\rangle\Big)\Big\}$.
  This defines the clause gadget for clause $\mathcal{C}_j = (\alpha_i \vee \beta_i \vee \gamma_i)$ as depicted in Fig. 7b.

Notice that $|V| \leq 1 + 4n + m = O(m+n)$ and $m_{\mathcal{A}} \leq 16n + 5m = O(m+n)$; therefore, the transformation is linearly bounded in time and space.

At this point, it is not difficult to see[2] that φ is `true` if and only if Γ_φ is dynamically-consistent. Indeed, we have already discussed the key points of the simulation of φ by means of Γ_φ.

To conclude, notice that any hyperarc $A \in \mathcal{A}$ of $\mathcal{H}_\varphi$ has weights $w_A(\cdot) \in \{-1, 0, 1\}$ and size $|A| \leq 3$. Since any hyperarc with three heads (tails) can be replaced by two hyperarcs each having at most two heads (tails), then General-CHyTN-DC remains PSPACE-hard even if $w_A(\cdot) \in \{-1, 0, 1\}$ and $|A| \leq 2$ for every multi-tail/head LTC $A \in \mathcal{A}$. Also notice that $w_a(\cdot) \in [-n-1, n+1] \cap \mathbf{Z}$ and $\ell_a \in \{p, \neg p \mid p \in P\} \cup \{\lambda\}$ holds for every weight w_a and label ℓ_a appearing in any *standard* LTC $a \in \mathcal{A}$. This concludes the proof. □

[2]Of course we will typeset all the details of this very last step of the proof, if the referees deem it appropriate.

Theorem 7 motivates the study of consistency problems on CHyTNs having either only multi-head or only multi-tail hyperarcs. Anyway, the notions and the algorithms that we are going to provide for multi-head CHyTNs admit a straightforward symmetric counterpart formulation in the setting of multi-tail CHyTNs. Also due this symmetry, from now on, we shall consider only multi-head CHyTNs.

4.1 ϵ-Dynamic-Consistency

In CSTNs/CHyTNs, decisions about the precise timing of actions are postponed until execution time, when informations meanwhile gathered at the observation nodes can be taken into account. However, the planner is allowed to factor in an outcome, and differentiate its strategy according to it, only strictly after the outcome has been observed (whence the strict inequality in Definition 14). Notice that this definition does not take into account the reaction time, which, in most applications, is non-negligible. In order to deliver algorithms that can also deal with the *reaction time* ϵ of the planner, we now employ a refined notion of dynamic-consistency. We stress that, in the forthcoming, our CHyTNs admit only multi-head hyperarcs.

Definition 19 (ϵ-dynamic-consistency)**.** *Given any CSTN/CHyTN $\langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P\rangle$ and any real number $\epsilon \in (0, +\infty)$, an execution strategy $\sigma \in \mathcal{S}_\Gamma$ is* ϵ-dynamic *if it satisfies all the H_ϵ-constraints, namely, for any two scenarios $s_1, s_2 \in \Sigma_P$ and any event $u \in V^+_{s_1,s_2}$, the execution strategy σ satisfies the following constraint, which is denoted $H_\epsilon(s_1; s_2; u)$:*

$$[\sigma(s_1)]_u \geq \min\Big(\{[\sigma(s_2)]_u\} \cup \{[\sigma(s_1)]_v + \epsilon \mid v \in \Delta(s_1; s_2)\}\Big)$$

We say that a CSTN/CHyTN Γ is ϵ-dynamically-consistent *if it admits $\sigma \in \mathcal{S}_\Gamma$ which is both viable and ϵ-dynamic. The problem of checking whether a given CSTN/CHyTN is ϵ-dynamically-consistent is named* CSTN-ϵ-DC/CHyTN-ϵ-DC.

It follows directly from Definition 19 that, whenever $\sigma \in S_\Gamma$ satisfies some $H_\epsilon(s_1; s_2; u)$, then σ satisfies $H_{\epsilon'}(s_1; s_2; u)$ for every $\epsilon' \in (0, \epsilon]$ as well. This proves the following lemma.

Lemma 5. *Let Γ be a CSTN or CHyTN. If Γ is ϵ-dynamically-consistent, for some $\epsilon > 0$, then Γ is ϵ'-dynamically-consistent for every $\epsilon' \in (0, \epsilon]$.*

Given any dynamically-consistent CSTN or CHyTN, we may ask for the maximum reaction time ϵ of the planner beyond which the network is no longer dynamically-consistent.

Definition 20 (Reaction time $\hat{\epsilon}$)**.** *Let Γ be a CSTN or CHyTN. Let $\hat{\epsilon} \triangleq \hat{\epsilon}(\Gamma)$ be the least upper bound of the set of all real numbers $\epsilon > 0$ such that Γ is ϵ-dynamically-consistent, i.e.,*

$$\hat{\epsilon} \triangleq \hat{\epsilon}(\Gamma) \triangleq \sup\{\epsilon > 0 \mid \Gamma \textit{ is } \epsilon\textit{-dynamically-consistent}\}.$$

Let us consider the (affinely) extended real numbers $\overline{\mathbf{R}} \triangleq \mathbf{R} \cup \{-\infty, \infty\}$, where *every* subset S of $\overline{\mathbf{R}}$ has an infimum and a supremum. In particular, recall that $\sup \emptyset = -\infty$ and, if S is unbounded above, then $\sup S = +\infty$.

If Γ is dynamically-consistent, then $\hat{\epsilon}(\Gamma)$ exists and $\hat{\epsilon}(\Gamma) \neq -\infty$ (i.e., the set on which we have taken the supremum is nonempty), as it is now proved in Lemma 6.

Lemma 6. *Let σ be a dynamic execution strategy for the CSTN/CHyTN Γ. Then, there exists a sufficiently small real number $\epsilon \in (0, +\infty)$ such that σ is ϵ-dynamic.*

Proof. Let $s_1, s_2 \in \Sigma_P$ be two scenarios and let us consider any event $u \in V^+_{s_1,s_2}$. Since σ is dynamic, then by Lemma 4 the following implication necessarily holds:

$$\Big(\bigwedge_{v \in \Delta(s_1;s_2)} [\sigma(s_1)]_u \leq [\sigma(s_1)]_v \Big) \Rightarrow [\sigma(s_1)]_u \geq [\sigma(s_2)]_u \qquad (*)$$

Notice that, w.r.t. Lemma 4, we have relaxed the equality $[\sigma(s_1)]_u = [\sigma(s_2)]_u$ in the implicand of (L4) by introducing the inequality $[\sigma(s_1)]_u \geq [\sigma(s_2)]_u$. At this point, we convert $(*)$ from implicative to disjunctive form, first by applying the rule of material implication[3], and then De Morgan's law[4]. From this, we see that the following disjunction must hold:

$$[\sigma(s_1)]_u \geq [\sigma(s_2)]_u \vee \bigvee_{v \in \Delta(s_1;s_2)} [\sigma(s_1)]_u > [\sigma(s_1)]_v \qquad (**)$$

Then, we argue that there exists a real number $\epsilon \in (0, +\infty)$ such that the following disjunction holds as well:

$$[\sigma(s_1)]_u \geq [\sigma(s_2)]_u \vee \bigvee_{v \in \Delta(s_1;s_2)} [\sigma(s_1)]_u \geq [\sigma(s_1)]_v + \epsilon.$$

In fact, since the disjunction $(**)$ necessarily holds, then one can pick the following real number ϵ:

$$\epsilon \triangleq \min_{\langle s_1,s_2,u \rangle \in \Sigma_P \times \Sigma_P \times V^+_{s_1,s_2}} \epsilon(s_1; s_2; u),$$

where the values $\epsilon(s_1; s_2; u) \in (0, +\infty)$ are defined as follows, for every $\langle s_1, s_2, u \rangle \in \Sigma_P \times \Sigma_P \times V^+_{s_1,s_2}$:

$$\epsilon(s_1; s_2; u) \triangleq \begin{cases} 1, \text{ if } [\sigma(s_1)]_u \geq [\sigma(s_2)]_u; \\ \min \Big\{ [\sigma(s_1)]_u - [\sigma(s_1)]_v \mid v \in \Delta(s_1; s_2),\ [\sigma(s_1)]_u > [\sigma(s_1)]_v \Big\}, \text{ otherwise.} \end{cases}$$

This implies that σ satisfies every H_ϵ-constraint of Γ, and thus that σ is ϵ-dynamic. □

Next, we prove a converse formulation of Lemma 6.

Lemma 7. *Let σ be an ϵ-dynamic execution strategy for a CSTN/CHyTN Γ, for some real number $\epsilon \in (0, +\infty)$.*

Then, σ is dynamic.

Proof. For the sake of contradiction, let us suppose that σ is not dynamic. Let F be the non-empty set of all the triplets $\langle u, s_1, s_2 \rangle \in V^+_{s_1,s_2} \times \Sigma_P \times \Sigma_P$, for which the implication (L4) does not hold. Then, $\langle u, s_1, s_2 \rangle \in F$ if and only if the following two properties hold:

1. $[\sigma(s_1)]_u \leq [\sigma(s_1)]_v$, for every $v \in \Delta(s_1; s_2)$;
2. $[\sigma(s_1)]_u \neq [\sigma(s_2)]_u$.

Let $\langle \hat{u}, \hat{s_1} \rangle$ be an event whose scheduling time is minimum and for which (1) and (2) hold, namely, let:

$$\langle \hat{u}, \hat{s_1} \rangle \triangleq \arg\min \Big\{ [\sigma(s_1)]_u \mid \exists s_2\, \langle u, s_1, s_2 \rangle \in F \Big\}.$$

[3]The rule of material implication: $\models p \Rightarrow q \iff \neg p \vee q$

[4]De Morgan's law: $\models \neg(p \wedge q) \iff \neg p \vee \neg q$

Since $\langle \hat{u}, \hat{s_1} \rangle$ is minimum in time, then $[\sigma(\hat{s_1})]_{\hat{u}} \leq [\sigma(s_2)]_{\hat{u}}$ for every $s_2 \in \Sigma_P$ such that $\langle \hat{u}, \hat{s_1}, s_2 \rangle \in F$; moreover, since $\langle \hat{u}, \hat{s_1}, s_2 \rangle \in F$, then $[\sigma(s_1)]_{\hat{u}} \neq [\sigma(s_2)]_{\hat{u}}$ holds by (2), so that $[\sigma(\hat{s_1})]_{\hat{u}} < [\sigma(s_2)]_{\hat{u}}$. At this point, recall that σ is ϵ-dynamic by hypothesis, hence $[\sigma(\hat{s_1})]_{\hat{u}} < [\sigma(s_2)]_{\hat{u}}$ implies that there exists $v \in \Delta(\hat{s_1}; s_2)$ such that:

$$[\sigma(\hat{s_1})]_{\hat{u}} \geq [\sigma(\hat{s_1})]_v + \epsilon > [\sigma(\hat{s_1})]_v,$$

but this inequality contradicts item (1) above. Indeed, $F = \emptyset$ and σ is thus dynamic. □

In Section 5, the following theorem is proved.

Theorem 8. *For any dynamically-consistent CSTN/CHyTN Γ, where V is the set of events and Σ_P is the set of scenarios, we have that $\hat{\epsilon}(\Gamma) \geq |\Sigma_P|^{-1}|V|^{-1}$.*

Notice that, in Definition 16, dynamic-consistency was defined by strict-inequality and equality constraints. However, by Theorem 9, dynamic-consistency can also be defined in terms of H_ϵ-constraints only (i.e., no strict-inequalities are required).

Theorem 9. *Let Γ be a CSTN/CHyTN. Let $\epsilon \triangleq |\Sigma_P|^{-1}|V|^{-1}$. Then, Γ is dynamically-consistent if and only if Γ is ϵ-dynamically-consistent.*

By Theorem 9, any algorithm for checking ϵ-dynamic-consistency can be used to check dynamic-consistency.

4.2 A (pseudo) Singly-Exponential Time Algorithm for CSTN-DC and CHyTN-DC

In this section, we present shall the first singly-exponential time algorithm for solving CSTN-DC and CHyTN-DC, also producing a dynamic execution strategy whenever the input CSTN/CHyTN is dynamically-consistent. Hereafter, let us denote $\mathbf{N}_0 \triangleq \mathbf{N} \setminus \{0\}$.

The main result of this paper is summarized in the following theorem, which is proven in this section.

Theorem 10. *The following two propositions hold.*

1. *There exists an $O(|\Sigma_P|^2|\mathcal{A}|^2 + |\Sigma_P|^3|\mathcal{A}||V||P| + |\Sigma_P|^4|V|^2|P|)WD$ time deterministic algorithm deciding CSTN-ϵ-DC/CHyTN-ϵ-DC on input $\langle \Gamma, \epsilon \rangle$, for any CSTN/CHyTN $\Gamma = \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P \rangle$ and any rational number $\epsilon = N/D$ where $N, D \in \mathbf{N}_0$. In particular, given any ϵ-dynamically-consistent CSTN/CHyTN Γ, the algorithm returns as output a viable and ϵ-dynamic execution strategy $\sigma : \Sigma_P \to \Phi_V$ for Γ.*
2. *There exists an $O(|\Sigma_P|^3|\mathcal{A}|^2|V| + |\Sigma_P|^4|\mathcal{A}||V|^2|P| + |\Sigma_P|^5|V|^3|P|)W$ time deterministic algorithm for checking CSTN-DC/CHyTN-ϵ-DC on any input $\Gamma = \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P \rangle$. In particular, given any dynamically-consistent CSTN/CHyTN Γ, it returns a viable and dynamic execution strategy $\sigma : \Sigma_P \to \Phi_V$ for Γ.*

Here, $W \triangleq \max_{a \in A} |w_a|$.

We now present the reduction from CHyTN-DC to HyTN-Consistency. Of course, since any CSTN is a CHyTN, the same argument reduces CSTN-DC to HyTN-Consistency. Firstly, we argue that any CSTN/CHyTN can be viewed as a succinct representation which can be expanded into an exponential sized STN/HyTN.

The *Expansion* of CSTNs is introduced below.

Definition 21 (Expansion $\langle V_\Gamma^{\text{Ex}}, \Lambda_\Gamma^{\text{Ex}} \rangle$)**.** *Let $\Gamma = \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P \rangle$ be a CSTN/CHyTN. Consider the distinct and disjoint STNs/HyTNs $\langle V_s, \mathcal{A}_s \rangle$, one for each scenario $s \in \Sigma_P$, defined as follows (where we adopt the notation $v_s \triangleq (v, s)$ for every $v \in V$ and $s \in \Sigma_P$):*

$$V_s \triangleq \{v_s \mid v \in V_s^+\}, \text{ and}$$

$$\mathcal{A}_s \triangleq \Big\{ \big(t_s, \underbrace{\{h_s^{(1)}, \ldots, h_s^{(k)}\}}_{\text{heads labeled with } s}, \underbrace{\langle w(h_s^{(1)}), \ldots, w(h_s^{(k)}) \rangle}_{\text{corresponding weights}} \big) \mid \exists k \, \big(\underbrace{t}_{\text{tail}}, \underbrace{\{h^{(1)}, \ldots, h^{(k)}\}}_{\text{heads}}, \underbrace{\langle w(h^{(1)}), \ldots, w(h^{(k)}) \rangle}_{\text{corresponding weights}} \big) \in \mathcal{A}_s^+ \Big\}.$$

(Of course, in the above notation, $k = 1$ when Γ is a CSTN, whereas $k \in \mathbf{N}_{\geq 1}$ when Γ is a CHyTN.)

Next, we define the expansion $\langle V_\Gamma^{Ex}, \Lambda_\Gamma^{Ex} \rangle$ *of Γ as follows:*

$$\langle V_\Gamma^{Ex}, \Lambda_\Gamma^{Ex} \rangle \triangleq \Big(\bigcup_{s \in \Sigma_P} V_s, \bigcup_{s \in \Sigma_P} \mathcal{A}_s \Big).$$

Notice that $V_{s_1} \cap V_{s_2} = \emptyset$ whenever $s_1 \neq s_2$ and that $\langle V_\Gamma^{\text{Ex}}, \Lambda_\Gamma^{\text{Ex}} \rangle$ is an STN/HyTN with at most $|V_\Gamma^{\text{Ex}}| \leq |\Sigma_P| \cdot |V|$ nodes and size at most $|\Lambda_\Gamma^{\text{Ex}}| \leq |\Sigma_P| \cdot |\mathcal{A}|$.

We now show that the expansion of a CSTN/CHyTN can be enriched with some (extra) multi-head hyperarcs in order to model ϵ-dynamic-consistency, by means of a particular HyTN which is denoted $\mathcal{H}_\epsilon(\Gamma)$.

Definition 22 (HyTN $\mathcal{H}_\epsilon(\Gamma)$)**.** *Let $\Gamma = \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P \rangle$ be a CSTN/CHyTN. Given any real number $\epsilon \in (0, +\infty)$, a corresponding HyTN denoted by $\mathcal{H}_\epsilon(\Gamma)$, is defined as follows:*

- *For every scenarios $s_1, s_2 \in \Sigma_P$ and for every event node $u \in V_{s_1,s_2}^+$, define a hyperarc $\alpha \triangleq \alpha_\epsilon(s_1; s_2; u)$ as follows (with the intention to model $H_\epsilon(s_1; s_2; u)$, see Def. 19):*

  $$\alpha_\epsilon(s_1; s_2; u) \triangleq \big(t_\alpha, H_\alpha, w_\alpha\big), \;\; \forall\, s_1, s_2 \in \Sigma_P \text{ and } u \in V_{s_1,s_2}^+.$$

  where:

  - $t_\alpha \triangleq u_{s_1}$ *is the tail of the (multi-head) hyperarc $\alpha_\epsilon(s_1; s_2; u)$;*
  - $H_\alpha \triangleq \{u_{s_2}\} \cup \Delta(s_1; s_2)$ *is the set of the heads of $\alpha_\epsilon(s_1; s_2; u)$;*
  - $w_\alpha(u_{s_2}) \triangleq 0$*; moreover, $w_\alpha(v) \triangleq -\epsilon$ for each $v \in \Delta(s_1; s_2)$.*

- *Consider the expansion $\langle V_\Gamma^{Ex}, \Lambda_\Gamma^{Ex} \rangle$ of Γ. Then, $\mathcal{H}_\epsilon(\Gamma)$ is defined as $\mathcal{H}_\epsilon(\Gamma) \triangleq \big(V_\Gamma^{Ex}, \mathcal{A}_{H_\epsilon}\big)$, where,*

  $$\mathcal{A}_{H_\epsilon} \triangleq \Lambda_\Gamma^{Ex} \cup \bigcup_{\substack{s_1, s_2 \in \Sigma_P \\ u \in V_{s_1,s_2}^+}} \alpha_\epsilon(s_1; s_2; u).$$

Notice that each $\alpha_\epsilon(s_1; s_2; u)$ has size $|\alpha_\epsilon(s_1; s_2; u)| = 1 + \Delta(s_1; s_2) \leq 1 + |P|$.

Here below, Algorithm 1 provides a pseudocode for constructing $\mathcal{H}_\epsilon(\Gamma)$.

Algorithm 1: `construct_`$\mathcal{H}(\Gamma, \epsilon)$

Input: a CSTN/CHyTN $\Gamma \triangleq \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P\rangle$, a rational $\epsilon > 0$

foreach $(s \in \Sigma_P)$ **do**
  $V_s \leftarrow \{v_s \mid v \in V_s^+\}$; // see Def. 21
  $A_s \leftarrow \{a_s \mid a \in A_s^+\}$; // see Def. 21
$V_\Gamma^{\text{Ex}} \leftarrow \cup_{s\in\Sigma_P} V_s$;
$\Lambda_\Gamma^{\text{Ex}} \leftarrow \cup_{s\in\Sigma_P} A_s$;
foreach $(s_1, s_2 \in \Sigma_P$ **and** $u \in V^+_{s_1,s_2})$ **do**
  $t_\alpha \leftarrow u_{s_1}$;
  $H_\alpha \leftarrow \{u_{s_2}\} \cup \Delta(s_1; s_2)$;
  $w_\alpha(u_{s_2}) \leftarrow 0$;
  foreach $v \in \Delta(s_1; s_2)$ **do**
    $w_\alpha(v_{s_1}) \leftarrow -\epsilon$;
  $\alpha_\epsilon(s_1; s_2; u) \leftarrow \big(t_\alpha, H_\alpha, w_\alpha\big)$;
$\mathcal{A}_{\mathcal{H}_\epsilon} \leftarrow \Lambda_\Gamma^{\text{Ex}} \cup \bigcup_{\substack{s_1,s_2\in\Sigma_P \\ u\in V^+_{s_1,s_2}}} \alpha_\epsilon(s_1; s_2; u)$;
$\mathcal{H}_\epsilon(\Gamma) \leftarrow \big(V_\Gamma^{\text{Ex}}, \mathcal{A}_{H_\epsilon}\big)$;
return $H_\epsilon(\Gamma)$;

An excerpt of the HyTN $\mathcal{H}_\epsilon(\Gamma)$ corresponding to the CSTN Γ_1 of Fig. 2a is depicted in Fig. 8.

The following theorem establishes the connection between dynamic-consistency of CSTNs/CHyTNs and consistency of HyTNs.

Theorem 11. *Given any CSTN/CHyTN $\Gamma = \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P\rangle$, there exists a sufficiently small real number $\epsilon \in (0, +\infty)$ such that the CSTN/CHyTN Γ is dynamically-consistent if and only if the HyTN $\mathcal{H}_\epsilon(\Gamma)$ is consistent.*

Moreover, the HyTN $\mathcal{H}_\epsilon(\Gamma)$ has at most $|V_{\mathcal{H}_\epsilon}| \leq |\Sigma_P|\,|V|$ nodes, $|\mathcal{A}_{\mathcal{H}_\epsilon}| = O(|\Sigma_P|\,|\mathcal{A}| + |\Sigma_P|^2|V|)$ hyperarcs, and it has size $m_{\mathcal{A}_{\mathcal{H}_\epsilon}} = O(|\Sigma_P|\,|\mathcal{A}| + |\Sigma_P|^2|V|\,|P|)$.

Proof. For any real number $\epsilon \in (0, +\infty)$, let $\mathcal{H}_\epsilon(\Gamma) = \langle V_\Gamma^{\text{Ex}}, \mathcal{A}_{H_\epsilon}\rangle$ be the HyTN of Definition 22.

(1) Firstly, we prove that, for any $\epsilon > 0$, $\mathcal{H}_\epsilon(\Gamma)$ is consistent if and only if Γ is ϵ-dynamically-consistent.

($\Rightarrow$) Given any feasible scheduling $\phi : V_\Gamma^{\text{Ex}} \to \mathbf{R}$ for the HyTN $\mathcal{H}_\epsilon(\Gamma)$, let $\sigma_\phi(s) \in \mathcal{S}_\Gamma$ be the execution strategy defined as follows:

$$[\sigma_\phi(s)]_v \triangleq \phi(v_s), \text{ for every } v_s \in V_\Gamma^{\text{E}}, \text{ where } v \in V \text{ and } s \in \Sigma_P.$$

Notice that each hyperarc $\alpha_\epsilon(s_1; s_2; u)$ is satisfied by ϕ if and only if the corresponding H_ϵ-constraint $H_\epsilon(s_1; s_2; u)$ is satisfied by σ_ϕ; moreover, recall that $\Lambda_\Gamma^{\text{Ex}} \subseteq \mathcal{A}_{H_\epsilon}$, and that $\Lambda_\Gamma^{\text{Ex}}$ contains all the *original* standard/hyper difference constraints of Γ (i.e., those induced by $\mathcal{A}$, by means of Def. 21). At this point, since ϕ is feasible for the HyTN $\mathcal{H}_\epsilon(\Gamma)$, then σ_ϕ must be viable and ϵ-dynamic for Γ (because it satisfies all the required constraints).

Therefore, Γ is ϵ-dynamically-consistent. □

($\Leftarrow$) Given any viable and ϵ-dynamic execution strategy $\sigma \in \mathcal{S}_\Gamma$, for some real number $\epsilon \in (0, +\infty)$, let $\phi_\sigma : V_\Gamma^{\text{Ex}} \to \mathbf{R}$ be the scheduling of the HyTN $\mathcal{H}_\epsilon(\Gamma)$ defined as follows:

$$\phi_\sigma(v_s) \triangleq [\sigma(s)]_v \text{ for every } v_s \in V_\Gamma^{\text{Ex}}, \text{ where } v \in V \text{ and } s \in \Sigma_P.$$

Also in this case, we have that $\Lambda_{\Gamma}^{\text{Ex}} \subseteq \mathcal{A}_{H_\epsilon}$, and a moment's reflection reveals that each hyperarc $\alpha_\epsilon(s_1; s_2; u)$ is satisfied by ϕ_σ if and only if $H_\epsilon(s_1; s_2; u)$ is satisfied by σ. At this point, since σ is viable and ϵ-dynamic for the CSTN/CHyTN Γ, then ϕ_σ must be feasible for $\mathcal{H}_\epsilon(\Gamma)$. Therefore, $\mathcal{H}_\epsilon(\Gamma)$ is consistent. □

This proves that, for any $\epsilon \in (0, +\infty)$, $\mathcal{H}_\epsilon(\Gamma)$ is consistent if and only if Γ is ϵ-dynamically-consistent.

(2) At this point, by composition with (1), Lemma 6 implies that there exists a sufficiently small real number $\epsilon \in (0, +\infty)$ such that Γ is dynamically-consistent if and only if $\mathcal{H}_\epsilon(\Gamma)$ is consistent.

(3) The size bounds follow directly from Definition 22.

This concludes the proof. □

At this point, we are in position to provide the pseudo-code for checking the CSTN-ϵ-DC/CHyTN-ϵ-DC of CSTNs/CHyTNs, it is given in Algorithm 2, here below:

Algorithm 2: `check_CSTN/CHyTN-`ϵ`-DC`(Γ, ϵ)

Input: a CSTN/CHyTN $\Gamma \triangleq \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P \rangle$, a rational number $\epsilon \triangleq N/D$, for $N, D \in \mathbf{N}_0$

1 $\mathcal{H}_\epsilon(\Gamma) \leftarrow$ `construct_`$\mathcal{H}(\Gamma, \epsilon)$; // ref. Algorithm 1
2 **foreach** $(A = \langle t_A, H_A, w_A \rangle \in \mathcal{A}_{\mathcal{H}_\epsilon(\Gamma)}$ **and** $h \in H_A)$ **do**
3 $\quad w_A(h) \leftarrow w_A(h)D$; // scale all weights of $\mathcal{H}_\epsilon(\Gamma)$, from $\mathbf{Q}$ to $\mathbf{Z}$
4 $\phi \leftarrow$ `check_HyTN-Consistency`$(\mathcal{H}_\epsilon(\Gamma))$; // ref. Thm 5
5 **if** (ϕ *is a feasible scheduling of* $\mathcal{H}_\epsilon(\Gamma)$) **then**
6 $\quad$ **foreach** (*event node* $v \in V_{\mathcal{H}_\epsilon(\Gamma)}$) **do**
7 $\quad\quad \phi(v) \leftarrow \phi(v)/D$; // re-scale back to size the scheduling time, from $\mathbf{Z}$ to $\mathbf{Q}$, w.r.t. ϵ
8 $\quad$ **return** $\langle$*YES*$, \phi\rangle$;
9 **else**
10 $\quad$ **return** *NO*;

whereas, the pseudo-code for checking CSTN-DC/CHyTN-DC is provided in Algorithm 3, here below:

Algorithm 3: `check_CSTN-DC/CHyTN-DC`(Γ)

Input: a CSTN/CHyTN $\Gamma \triangleq \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P \rangle$

1 $\hat{\epsilon} \leftarrow |\Sigma_P|^{-1}|V|^{-1}$; // ref. Thm. 9
2 **return** *check_CSTN/CHyTN-*ϵ*-DC*$(\Gamma, \hat{\epsilon})$;

The latter (Algorithm 3) invokes the former (Algorithm 2), and they goes as follows.

Description of the Algorithm Firstly, Algorithm 3 computes a sufficiently small *rational* number $\epsilon \in (0, +\infty) \cap \mathbf{Q}$, choosing it by resorting to Theorem 9, i.e., $\hat{\epsilon} \triangleq |\Sigma_P|^{-1}|V|^{-1}$ (at line 1 of Algorithm 3). Secondly, it invokes Algorithm 2 on input $\Gamma, \hat{\epsilon}$. At this point, Algorithm 2 firstly constructs $\mathcal{H}_{\hat{\epsilon}}(\Gamma)$ (at line 1 of Algorithm 2) by invoking Algorithm 1, and then it scales *every* hyperarc's weight, appearing in $\mathcal{H}_{\hat{\epsilon}}(\Gamma)$, from $\mathbf{Q}$ to $\mathbf{Z}$ (at lines 2-3). This is done by multiplying each weight by a factor D (line 3), where $D \in \mathbf{N}_0$ is the denominator of $\hat{\epsilon}$ (i.e., $D = |\Sigma_P||V|$). Thirdly, $\mathcal{H}_{\hat{\epsilon}}(\Gamma)$ is solved with the HyTN-Consistency-Checking algorithm underlying Theorem 5 (at line 4), i.e., within the underlying algorithmic engine, an instance of the HyTN-Consistency problem gets solved by reducing it to the problem of determining winning regions in a carefully constructed MPGs (see [11, 12] for the details of such construction). At this point, if the HyTN-Consistency algorithm outputs YES, together with a feasible scheduling ϕ of $\mathcal{H}_{\hat{\epsilon}}(\Gamma)$, then the time values of ϕ are scaled back to size w.r.t. $\hat{\epsilon}$, and then $\langle$YES$, \phi\rangle$ is returned as output (lines 5-8); otherwise, the output is simply NO (at line 10). Still, notice that, thanks to the Item 3 of Theorem 5, we could return a negative certificate as well, because it is well characterized in terms of generalized negative cycles (see Definition 5).

Remark 2. *The same algorithm, with essentially the same upper bound on its running time and space, work also in case we allow for arbitrary boolean formulae as labels, rather than just conjunctions.*

Remark 3. *We remark that the HyTN/MPG algorithm that is at the heart of our approach requires integral weights (i.e., it requires that $w(u,v) \in \mathbf{Z}$ for every $(u,v) \in A$), and, somehow, we could not play it differently (see [11, 12] for a discussion). Moreover, the algorithm always computes an integral solution to HyTN/MPGs and, therefore, it always computes* rational *feasible schedules for the CSTNs/CHyTNs given in input. As such, it seems to us that this "requirement" actually turns out to be a plus, in practice. To conclude, it is indeed the integrality assumption that allows us to analyze the algorithm quantitatively, also presenting a sharp lower bounding analysis on the critical value of the reaction time $\hat{\epsilon}$, where the CSTN/CHyTN transits from being, to not being, dynamically consistent. We believe that these issues deserve much attention, going into them required a "discrete" approach to the notion of numbers.*

Now, the correctness and the time complexity of Algorithm 3 is analyzed. To begin, notice that some of the temporal constraints introduced during the reduction step depends on a sufficiently small parameter $\hat{\epsilon} \in (0, +\infty) \cap \mathbf{Q}$, whose magnitude turns out to depend on the size of the input CSTN/CHyTN. It is now proved that the time complexity of the algorithm depends multiplicatively on D, provided that $\hat{\epsilon} = N/D$ for some $N, D \in \mathbf{N}_0$. In Section 5 we will present a sharp lower bounding analysis on $\hat{\epsilon}$, from which the (pseudo) singly-exponential time bound follows as corollary.

So, let us assume for a moment that line 1 of Algorithm 3 is valid, we shall prove this in Theorem 8.

Then, as a corollary of Theorem 11, we have that Algorithm 3 correctly decides CSTN-DC.

The most time expensive step of the algorithm is clearly line 4 of Algorithm 2, which resorts to Theorem 5 in order to solve an instance of HyTN-Consistency. From Theorem 11 we have an upper bound on the size of $\mathcal{H}_\epsilon(\Gamma)$, while Theorem 5 gives us a pseudo-polynomial upper bound for the computation time. Also, recall that we scale weights by a factor D at lines 2-3 of Algorithm 2, where $\hat{\epsilon} = N/D$ for some $N, D \in \mathbf{N}_0$. Thus, by composition, Algorithm 3 decides CSTN-DC/CHyTN-DC in a time $T^{A_3}_{|\Gamma|}$ which is bounded as follows, where $W \triangleq \max_{a \in A} |w_a|$:

$$T^{A_3}_{|\Gamma|} = O\Big(\big(|V_{H_\epsilon(\Gamma)}| + |\mathcal{A}_{\mathcal{H}_\epsilon(\Gamma)}|\big) m_{\mathcal{A}_{\mathcal{H}_\epsilon(\Gamma)}} \Big) WD.$$

Whence, taking into account the upper bound on the size of $\mathcal{H}_\epsilon(\Gamma)$ give by Theorem 11, the following holds:

$$T^{A_3}_{|\Gamma|} = O(|\Sigma_P|^2|A|^2 + |\Sigma_P|^3|A||V||P| + |\Sigma_P|^4|V|^2|P|)WD.$$

By Theorem 8, it is sufficient to check ϵ-dynamic-consistency for $\hat{\epsilon} = |\Sigma_P|^{-1}|V|^{-1}$. Therefore, an

$$O(|\Sigma_P|^3|V||A|^2 + |\Sigma_P|^4|A||V|^2|P| + |\Sigma_P|^5|V|^3|P|)W$$

worst-case time bound follows for Algorithm 3.

Since $|\Sigma_P| \leq 2^{|P|}$, the (pseudo) singly-exponential time bound follows. This proves Theorem 10.

5 Bounding Analysis on the Reaction Time $\hat{\epsilon}$

In this section we present an asymptotically sharp lower bound for $\hat{\epsilon}(\Gamma)$, that is the critical value of reaction time where the CSTN/CHyTN transits from being, to not being, dynamically-consistent. The proof technique introduced in this analysis is applicable more in general, when dealing with linear difference constraints which include strict inequalities. Moreover, this bound implies that Algorithm 3 is a (pseudo) singly-exponential time algorithm for solving CSTN-DC/CHyTN-DC.

To begin, we are going to provide a proof of Theorem 8, but, for this, let us first introduce some further notation.

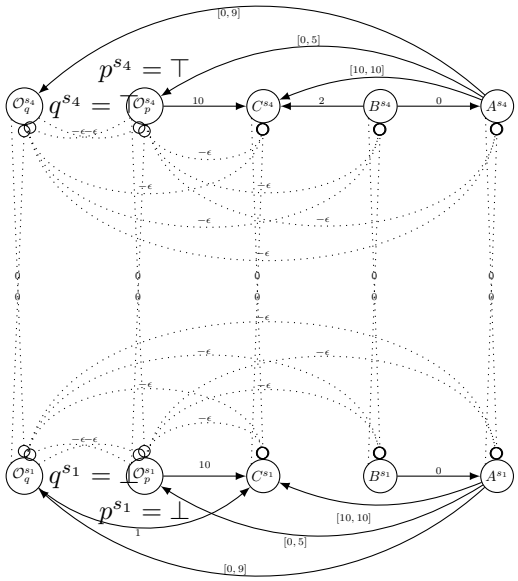

Figure 8: An excerpt of the HyTN $\mathcal{H}_\epsilon(\Gamma)$ corresponding to the CSTN Γ of Fig. 2a, in which two scenarios, s_4 and s_1, are considered and the corresponding hyper-constraints $H(s_4; s_1; u)$ are depicted as dotted hyperarcs.

Let $\Gamma \triangleq \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P\rangle$ be a dynamically-consistent CSTN/CHyTN. By Theorem 11, there exists $\epsilon > 0$ such that $\mathcal{H}_\epsilon(\Gamma)$ is consistent. Then, let $\phi : V_\Gamma^{\text{Ex}} \to \mathbf{R}$ be a feasible scheduling for $\mathcal{H}_\epsilon(\Gamma)$. For any hyperarc $A = \langle t_A, H_A, w_A\rangle \in \mathcal{A}_{\mathcal{H}_\epsilon}$, define a standard arc a_A as follows:

$$a_A \triangleq \langle t_A, \hat{h}, w_A(\hat{h})\rangle, \text{ where } \hat{h} \triangleq \arg\min_{h \in H_A} \big(\phi(h) - w_A(h)\big).$$

Then, notice that the network $T_\epsilon^\phi(\Gamma) \triangleq \langle V_\Gamma^{\text{Ex}}, \bigcup_{A\in\mathcal{A}_{\mathcal{H}_\epsilon}} a_A\rangle$ is always an STN. Moreover, a moment's reflection revelas that, as we picked the argmin's head $\hat{h}$ (in the definition above), ϕ is a feasible scheduling for the STN $T_\epsilon^\phi(\Gamma)$.

At this point, assuming $v \in V_\Gamma^{\text{Ex}}$, let us consider the *fractional part* r_v of ϕ_v, i.e.,

$$r_v \triangleq \phi_v - \lfloor \phi_v \rfloor.$$

Then, let $R \triangleq \{r_v\}_{v\in V_\Gamma^{\text{Ex}}}$ be the set of all the fractional parts induced by V_Γ^{Ex}. Sort R by the common ordering on $\mathbf{R}$ and assume that $S \triangleq \{r_1, \ldots, r_k\}$ is the resulting ordered set (without repetitions), i.e., $|S| = k$, $S = R$, $r_1 < \ldots < r_k$. Now, let $\mathsf{pos}(v)$ be the index position $1 \le \mathsf{pos}(v) \le k$ such that:

$$r_{\mathsf{pos}(v)} = r_v.$$

Then, we define a new fractional part as follows:

$$r'_v \triangleq \frac{\mathsf{pos}(v) - 1}{|\Sigma_P||V|} \tag{NFP}$$

also, we define a new scheduling function as follows:

$$\phi'_v \triangleq \lfloor \phi_v \rfloor + r'_v \tag{NSF}$$

Then, notice that the following holds.

Remark 4. *Notice that (NFP) doesn't alter the ordering relation among the fractional parts, i.e.,*

$$r'_u < r'_v \iff r_u < r_v, \text{ for any } u, v \in V_\Gamma^{Ex},$$

moreover, observe that (NSF) doesn't change the value of any integral part, i.e.,

$$\lfloor \phi'_u \rfloor = \lfloor \phi_u \rfloor, \text{ for any } u \in V_\Gamma^{Ex}.$$

We are now in the position to prove Theorem 8.

Proof of Theorem 8. Let $\Gamma \triangleq \langle V, \mathcal{A}, L, \mathcal{O}, \mathcal{O}V, P \rangle$ be dynamically-consistent. By Theorem 11 there exists $\epsilon' > 0$ such that $\mathcal{H}_{\epsilon'}(\Gamma)$ is consistent and admits some *feasible* scheduling $\phi : V_\Gamma^{\mathrm{Ex}} \to \mathbf{R}$. As mentioned, ϕ is feasible for the STN $T^\phi_{\epsilon'}(\Gamma)$. Now, let $\hat{\epsilon} \triangleq |\Sigma_P|^{-1}|V|^{-1}$. Moreover, let $T^\phi_{\hat{\epsilon}}(\Gamma)$ be the STN obtained from $T^\phi_{\epsilon'}(\Gamma)$ simply by replacing, in the weights of the arcs, each weight $-\epsilon'$ by $-\hat{\epsilon}$. We argue that ϕ' (as defined in (NSF) w.r.t. ϕ, V, Σ_P), is a feasible scheduling for the STN $T^\phi_{\hat{\epsilon}}(\Gamma)$. Indeed, every constraint of $T^\phi_{\hat{\epsilon}}(\Gamma)$ has form $\phi_v - \phi_u \leq w$, for $w \in \mathbf{Z}$ or $w = -\hat{\epsilon}$.

- Consider the case $w \in \mathbf{Z}$. Notice that $\phi_v - \phi_u \leq w$ holds because ϕ is feasible for $T^\phi_{\epsilon'}(\Gamma)$. Then, it is not difficult to see that $\phi'_v - \phi'_u \leq w$ holds as well, because of Remark 4.
- Consider the case $w = -\hat{\epsilon}$. Notice that $\phi_v - \phi_u \leq -\epsilon'$ holds because ϕ is feasible for the STN $T^\phi_{\epsilon'}(\Gamma)$. Then, notice the following implication always holds,

  $$\phi_v - \phi_u \leq -\epsilon' \Longrightarrow \phi_v \neq \phi_u.$$

  Hence, again by Remark 4, we can conclude that $\phi'_v \neq \phi'_u$. At this point, we observe that the temporal distance between ϕ'_u and ϕ'_v is, therefore, at least $\hat{\epsilon}$ by definition of (NSF) and (NFP), i.e.,

  $$\phi'_u - \phi'_v \geq |\Sigma_P|^{-1}|V|^{-1} = \hat{\epsilon}.$$

  That is to say, $\phi'_v - \phi'_u \leq -\hat{\epsilon}$.

This proves that ϕ' is a feasible scheduling also for the STN $T^\phi_{\hat{\epsilon}}(\Gamma)$. Since $T^\phi_{\hat{\epsilon}}(\Gamma)$ is thus consistent, then, a moment's reflection reveals that $\mathcal{H}_{\hat{\epsilon}}(\Gamma)$ is consistent as well thanks to the same scheduling function ϕ'.

Therefore, by Theorem 11, the CSTN/CHyTN Γ is $\hat{\epsilon}$-dynamically-consistent, for $\hat{\epsilon} \triangleq |\Sigma_P|^{-1}|V|^{-1}$. □

At this point, a natural question is whether the lower bound given by Theorem 8 can be improved up to $\hat{\epsilon}(\Gamma) = \Omega(|V|^{-1})$. In turn, this would improve the time complexity for Algorithm 3 by a factor $|\Sigma_P|$. However, the following theorem shows that this is not the case, by exhibiting a CSTN for which $\hat{\epsilon}(\Gamma) = 2^{-\Omega(|P|)}$. This proves that the lower bound given by Theorem 8 is (almost) asymptotically sharp.

Theorem 12. *For each $n \in \mathbf{N}_0$ there exists a CSTN Γ^n such that:*

$$\hat{\epsilon}(\Gamma^n) < 2^{-n+1} = 2^{-|P^n|/3+1},$$

where P^n is the set of boolean variables of Γ^n.

Proof. For each $n \in \mathbf{N}_0$, define a CSTN $\Gamma^n \triangleq \langle V^n, A^n, L^n, \mathcal{O}^n, \mathcal{O}V^n, P^n \rangle$ as follows. See Fig. 9 for a clarifying illustration.

- $V^n \triangleq \{X_i, Y_i, Z_i \mid 1 \leq i \leq n\}$;

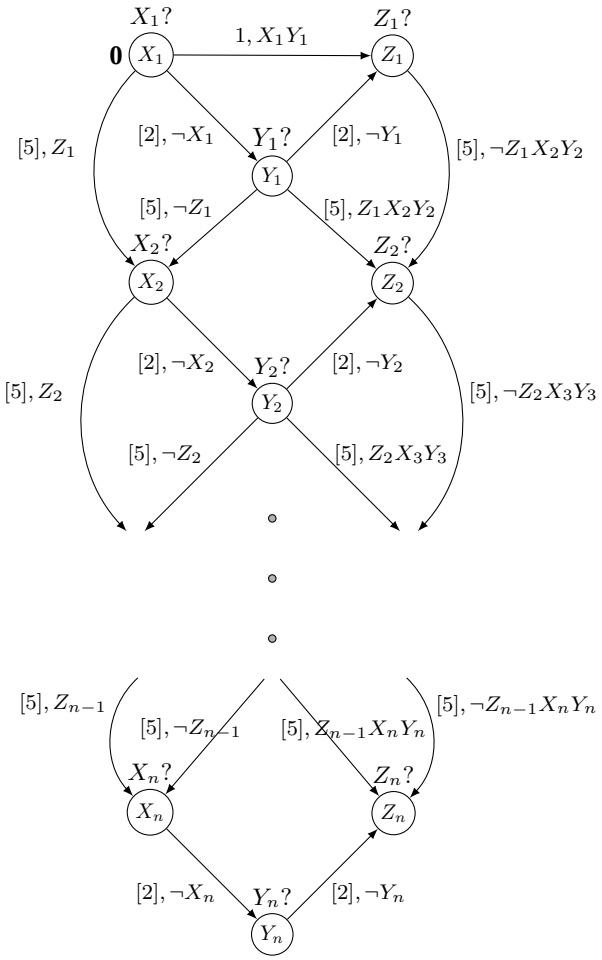

Figure 9: A CSTN Γ^n such that $\hat{\epsilon}(\Gamma^n) = 2^{-\Omega(|P^n|)}$.

- $A^n \triangleq B \cup \bigcup_{i=1}^{n} C_i \cup \bigcup_{i=1}^{n-1} D_i \cup E$
  where:
  - $B \triangleq \{\langle X_1 - v \leq 0, \lambda\rangle \mid v \in V^n\} \cup \{\langle Z_1 - X_1 \leq 1, X_1 \wedge Y_1\rangle\}$;
  - $C_i \triangleq \{\langle Y_i - X_i \leq 2, \neg X_i\rangle, \langle X_i - Y_i \leq -2, \neg X_i\rangle, \langle Z_i - Y_i \leq 2, \neg Y_i\rangle, \langle Y_i - Z_i \leq -2, \neg Y_i\rangle\}$;
  - $D_i \triangleq \{\langle X_{i+1} - X_i \leq 5, Z_i\rangle, \langle X_i - X_{i+1} \leq -5, Z_i\rangle, \langle X_{i+1} - Y_i\rangle \leq 5, \neg Z_i\rangle, \langle Y_i - X_{i+1} \leq -5, \neg Z_i\rangle, \langle Z_{i+1} - Y_i \leq 5, Z_i \wedge X_{i+1} \wedge Y_{i+1}\rangle, \langle Y_i - Z_{i+1} \leq -5, Z_i \wedge X_{i+1} \wedge Y_{i+1}\rangle, \langle Z_{i+1} - Z_i \leq 5, \neg Z_i \wedge X_{i+1} \wedge Y_{i+1}\rangle, \langle Z_i - Z_{i+1} \leq -5, \neg Z_i \wedge X_{i+1} \wedge Y_{i+1}\rangle\}$;
  - $E \triangleq \{\langle Y_n - X_n \leq 2, \neg X_n\rangle, \langle X_n - Y_n \leq -2, \neg X_n\rangle, \langle Z_n - Y_n \leq 2, \neg Y_n\rangle, \langle Y_n - Z_n \leq -2, \neg Y_n\rangle\}$;
- $L^n(v) \triangleq \lambda$ for every $v \in V^n$; $\mathcal{O}V^n \triangleq V^n$; $\mathcal{O}^n(v) \triangleq v$ for every $v \in \mathcal{O}V^n$; $P^n \triangleq V^n$.

We exhibit a viable and dynamic execution strategy $\sigma_n : \Sigma_{P^n} \to \Phi_{V^n}$ for Γ^n.

Let $\{\delta_i\}_{i=1}^n$ and $\{\Delta_i\}_{i=1}^n$ be two real valued sequences s.t.:

$$(1)\ \Delta_1 \triangleq 1; (2)\ 0 < \delta_i < \Delta_i; (3)\ \Delta_i \triangleq \min(\delta_{i-1}, \Delta_{i-1} - \delta_{i-1}).$$

Then, the following also holds for every $1 \leq i \leq n$:

$$(4)\ 0 < \Delta_i \leq 2^{-i+1},$$

where the equality holds if and only if $\delta_i = \Delta_i/2$.

In what follows, provided that $s \in \Sigma_P$ and $\ell \in P^*$, we will denote $\mathbb{1}_{s(\ell)} \triangleq 1$ if $s(\ell) = \top$ and $\mathbb{1}_{s(\ell)} \triangleq 0$ if $s(\ell) = \bot$.

We are ready to define $\sigma_n(s)$ for any $s \in \Sigma_P$:

- $[\sigma_n(s)]_{X_1} \triangleq 0$;
- $[\sigma_n(s)]_{Y_1} \triangleq \delta_1 \mathbb{1}_{s(X_1)} + 2\mathbb{1}_{s(\neg X_1)}$;
- $[\sigma_n(s)]_{Z_1} \triangleq \mathbb{1}_{s(X_1 \wedge Y_1)} + (2 + [\sigma_n(s)]_{Y_1})\mathbb{1}_{s(\neg X_1 \vee \neg Y_1)}$;
- $[\sigma_n(s)]_{X_i} \triangleq 5 + [\sigma_n(s)]_{X_{i-1}} \mathbb{1}_{s(Z_{i-1})} +$
  $+ [\sigma_n(s)]_{Y_{i-1}} \mathbb{1}_{s(\neg Z_{i-1})}$, for any $2 \leq i \leq n$;
- $[\sigma_n(s)]_{Y_i} \triangleq [\sigma_n(s)]_{X_i} + \delta_i \mathbb{1}_{s(X_i)} + 2\mathbb{1}_{s(\neg X_i)}$, for any $2 \leq i \leq n$;
- $[\sigma_n(s)]_{Z_i} \triangleq \big(5 + [\sigma_n(s)]_{Y_{i-1}} \mathbb{1}_{s(Z_{i-1})} +$
  $+ [\sigma_n(s)]_{Z_{i-1}} \mathbb{1}_{s(\neg Z_{i-1})}\big) \mathbb{1}_{s(X_i \wedge Y_i)} +$
  $+ (2 + [\sigma_n(s)]_{Y_i}) \mathbb{1}_{s(\neg X_i \vee \neg Y_i)}$, for any $2 \leq i \leq n$;

Let us prove, by induction on $n \geq 1$, that σ_n is viable and dynamic for Γ^n.

- *Base case*. Let $n = 1$. Then, by construction, σ_1 leads to the scheduling depicted in Figure 10. This shows that σ_1 is viable and dynamic for Γ^1.

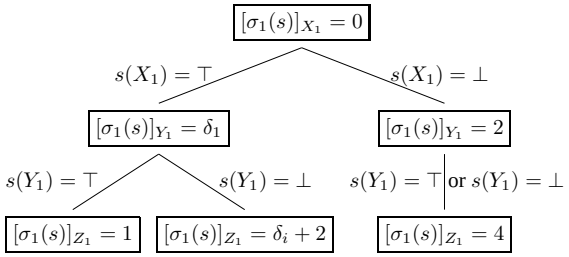

Figure 10: A viable and dynamic execution strategy for the base case $n = 1$.

- *Inductive step*. Let us assume that σ_{n-1} is viable and dynamic for Γ^{n-1}. Then, by construction, $[\sigma_n(s)]_v = [\sigma_{n-1}(s)]_v$ for every $s \in \Sigma_P$ and for every $v \in V_{n-1}$. Hence, by induction hypothesis, σ_n is viable and dynamic on V_{n-1}. Moreover, σ_n gives rise by construction to the scheduling depicted in Figure 11 and Figure 12. This shows that σ_n is viable and dynamic even for $V_n \setminus V_{n-1}$. Thus, σ_n is viable and dynamic for Γ^n, i.e., Γ^n is dynamically-consistent.

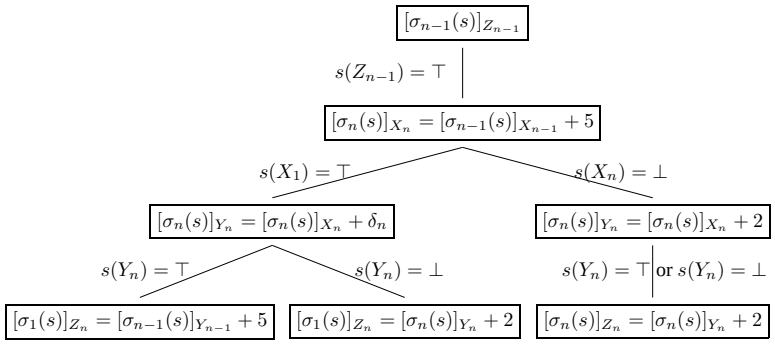
$[\sigma_{n-1}(s)]_{Z_{n-1}}$
$s(Z_{n-1}) = \top$
$[\sigma_n(s)]_{X_n} = [\sigma_{n-1}(s)]_{X_{n-1}} + 5$
$s(X_1) = \top$
$s(X_n) = \bot$
$[\sigma_n(s)]_{Y_n} = [\sigma_n(s)]_{X_n} + \delta_n$
$[\sigma_n(s)]_{Y_n} = [\sigma_n(s)]_{X_n} + 2$
$s(Y_n) = \top$
$s(Y_n) = \bot$
$s(Y_n) = \top$ or $s(Y_n) = \bot$
$[\sigma_1(s)]_{Z_n} = [\sigma_{n-1}(s)]_{Y_{n-1}} + 5$
$[\sigma_1(s)]_{Z_n} = [\sigma_n(s)]_{Y_n} + 2$
$[\sigma_n(s)]_{Z_n} = [\sigma_n(s)]_{Y_n} + 2$

Figure 11: A viable and dynamic execution strategy for the inductive step $n-1 \rightsquigarrow n$ when $s(Z_{n-1}) = \top$.

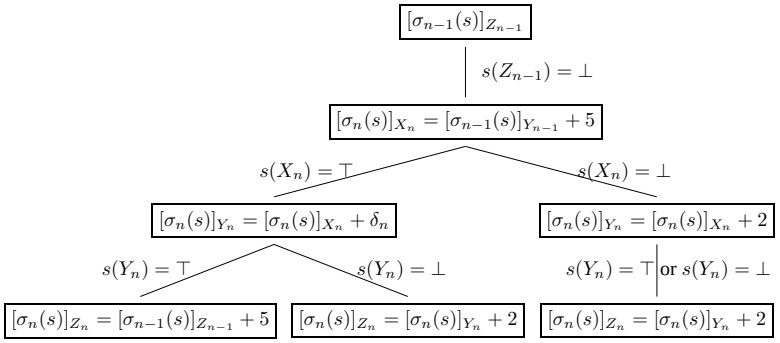
$[\sigma_{n-1}(s)]_{Z_{n-1}}$
$s(Z_{n-1}) = \bot$
$[\sigma_n(s)]_{X_n} = [\sigma_{n-1}(s)]_{Y_{n-1}} + 5$
$s(X_n) = \top$
$s(X_n) = \bot$
$[\sigma_n(s)]_{Y_n} = [\sigma_n(s)]_{X_n} + \delta_n$
$[\sigma_n(s)]_{Y_n} = [\sigma_n(s)]_{X_n} + 2$
$s(Y_n) = \top$
$s(Y_n) = \bot$
$s(Y_n) = \top$ or $s(Y_n) = \bot$
$[\sigma_n(s)]_{Z_n} = [\sigma_{n-1}(s)]_{Z_{n-1}} + 5$
$[\sigma_n(s)]_{Z_n} = [\sigma_n(s)]_{Y_n} + 2$
$[\sigma_n(s)]_{Z_n} = [\sigma_n(s)]_{Y_n} + 2$

Figure 12: A viable and dynamic execution strategy for the inductive step $n-1 \rightsquigarrow n$ when $s(Z_{n-1}) = \bot$.

Next, we show that $\hat{\epsilon}(\Gamma^n) < 2^{-n+1} = 2^{-|P^n|/3+1}$ for every $n \geq 1$. Let us consider the following scenario $\hat{s}$ for $1 \leq i \leq n$:

$$\hat{s}(X_i) \triangleq \hat{s}(Y_i) \triangleq \top; \quad \hat{s}(Z_i) \triangleq \begin{cases} \top, & \text{if } \delta_i \leq \Delta_i/2 \\ \bot, & \text{if } \delta_i > \Delta_i/2 \end{cases}.$$

We assume that σ is an execution strategy for Γ^n and study necessary conditions to ensure that σ is viable and dynamic, provided that the observations follow scenario $\hat{s}$. First, σ must schedule X_1 at time $[\sigma(\hat{s})]_{X_1} = 0$. Then, since $\hat{s}(X_1) = \top$, we must have $0 < [\sigma(\hat{s})]_{Y_1} < 1$, because of the constraint $(Z_1 - X_1 \leq 1, X_1 \wedge Y_1)$. Stated otherwise, it is necessary that:

$$0 < [\sigma(\hat{s})]_{Y_1} - [\sigma(\hat{s})]_{X_1} < \Delta_1.$$

After that, since $\hat{s}(Y_1) = \top$, then σ must schedule Z_1 at time $[\sigma(\hat{s})]_{Z_1} = 1 = \Delta_1$. A moment's reflection reveals that almost identical necessary conditions now recur for X_2, Y_2, Z_2, with the crucial variation that it will be necessary to require: $0 < [\sigma(\hat{s})]_{Y_2} < \Delta_2$. Indeed, proceeding inductively, it will be necessary that for every $1 \leq i \leq n$ and every $n \in \mathbf{N}_0$:

$$0 < [\sigma(\hat{s})]_{Y_i} - [\sigma(\hat{s})]_{X_i} < \Delta_i.$$

As already observed in (4), we have $0 < \Delta_n \leq 2^{-n+1}$. Thus, any viable and dynamic execution strategy σ for Γ^n must satisfy:

$$0 < [\sigma(\hat{s})]_{Y_n} - [\sigma(\hat{s})]_{X_n} < \frac{1}{2^{n-1}} = \frac{1}{2^{|P^n|/3-1}}.$$

Therefore, once the planner has observed the outcome $\hat{s}(X_n) = \top$ from the observation event X_n, then he must react by scheduling Y_n within time $2^{-n+1} = 2^{-|P^n|/3+1}$ in the future w.r.t. $[\sigma(\hat{s})]_{X_n}$.

Then, $\hat{\epsilon}(\Gamma^n) < 2^{-n+1} = 2^{-|P^n|/3+1}$ holds for every $n \geq 1$. □

6 Related Works

This section discusses of some alternative approaches offered by the current literature. Recall that the article of [28] has been discussed already in the introduction.

The work of [7] provided the first sound-and-complete algorithm for checking the dynamic-controllability of CSTNs with Uncertainty (CSTNU) and thus it can be employed for checking the dynamic-consistency of CSTNs, as a special case. The algorithm reduces the DC-Checking of CSTNUs to the problem of solving Timed Game Automata (TGA). Nevertheless, no worst-case bound on the time complexity of the procedure was provided in [7]. Still, one may observe that solving TGAs is a problem of much higher complexity than solving MPGs, compare the following known facts: solving 1-player TGAs is PSPACE-complete and solving 2-player TGAs is EXP-complete; on the contrary, the problem of determining MPGs lies in NP ∩ coNP and it is currently an open problem to prove whether it is in P. Indeed, the algorithm presented in [7] is not singly-exponential time bounded.

Next, a sound algorithm for checking the dynamic-controllability of CSTNUs was given by Combi, Hunsberger, Posenato in [8]. However, it was not shown to be complete. To the best of our knowledge, it is currently open whether or not it can be extended in order to prove completeness for the CSTNU model.

More about the particular CSTN model, [22] presented, at the same conference in which it appeared the preliminary version of this work, a sound-and-complete DC-checking algorithm for CSTNs. It is based on the propagation of constraints labeled by propositions. However, to the best of our knowledge, the worst-case complexity of the algorithm is currently unsettled. Also notice that the algorithm in [22] works on CSTNs, and doesn't take into account the CHyTN model. Indeed, we strongly believe that our approach (based on tractable games and reaction time analysis of $\hat{\epsilon}$) and the approach of [22] (based on the propagation of constraints) can benefit a lot from each other. Further studies are needed to understand where and what are the points and opportunities for a contact. Also on this route, we believe, there's still plenty that can be done.

7 Conclusion

In this work we introduced the *Conditional Hyper Temporal Network (CHyTN)* model, which is a natural extension and generalization of both the CSTN and the HyTN model, obtained by "melting" them together. We proved that deciding whether a given CSTN or CHyTN is dynamically consistent is coNP-hard; and that deciding whether a given CHyTN is dynamically consistent is PSPACE-hard, provided that the input instances are allowed to include both multi-head and multi-tail hyperarcs. In light of this, we thus focused on CHyTNs that allow only multi-head or only multi-tail hyperarcs, and we offered the first deterministic (pseudo) singly-exponential time algorithm for the problem of checking the dynamic-consistency of such CHyTNs, also producing a dynamic execution strategy whenever the input CHyTN is dynamically consistent. As a byproduct, this provides the first sound-and-complete (pseudo) singly-exponential time algorithm for checking the dynamic-consistency of CSTNs. The algorithm is based on a novel connection between CSTNs/CHyTNs and Mean Payoff Games. The presentation of the connection between CSTNs/CHyTNs and MPGs was mediated by the HyTN model. The algorithm actually manages a few more general variants of the problem; e.g., those where labels are not required to be conjunctions. To summarize, at the heart of the algorithm a reduction to MPGs is mediated by the HyTN model. The CSTN/CHyTN is dynamically consistent if and only if the corresponding MPG is everywhere won, and a dynamic execution strategy can be conveniently read out by an everywhere winning positional strategy. The size of this MPG is at most

polynomial in the number of the possible scenarios; as such, the term at the corresponding time complexity's exponent is linear, at worst, in the number of the observation events. The same holds for the running time of the resulting algorithm. Finally, in order to analyze the algorithm, we introduced a refined notion of dynamic-consistency, named ϵ-dynamic-consistency, also presenting a sharp lower bounding analysis on the critical value of the reaction time $\hat{\varepsilon}$ where a CSTN/CHyTN transits from being, to not being, dynamically consistent.

In future works we would like to settle the exact computational complexity of CSTN-DC/CHyTN-DC, as well as to extend our approach in order to check the dynamic-controllability of CSTN with Uncertainty [21]. Finally, an extensive experimental evaluation which takes good account of all the possible optimization and heuristics is planned.

Acknowledgments This work was supported by the *Department of Computer Science, University of Verona, Verona, Italy* under *Ph.D.* grant "Computational Mathematics and Biology" on a co-tutelle agreement with *Laboratoire d'Informatique Gaspard-Monge (LIGM), Université Paris-Est, Marne-la-Vallée, Paris, France*.

References

[1] X. Allamigeon, P. Benchimol, and S. Gaubert. The tropical shadow-vertex algorithm solves mean payoff games in polynomial time on average. In *ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, pages 89–100, 2014.

[2] Richard Bellman. On a routing problem. *Quarterly of Applied Mathematics*, 16(1):87–90, 1958.

[3] Claudio Bettini, Xiaoyang Sean Wang, and Sushil Jajodia. Temporal reasoning in workflow systems. *Dist. & Paral. Data.*, 11(3):269–306, 2002.

[4] L. Brim and J. Chaloupka. Using strategy improvement to stay alive. *Int. J. Found. Comput. Sci.*, 23(3):585–608, 2012.

[5] L. Brim, J. Chaloupka, L. Doyen, R. Gentilini, and J. F. Raskin. Faster algorithms for mean-payoff games. *Formal Methods in System Design*, 38(2):97–118, 2011.

[6] K. Chatterjee, M. Henzinger, S. Krinninger, and D. Nanongkai. Polynomial-time algorithms for energy games with special weight structures. *Algorithmica*, 70(3):457–492, 2014.

[7] A. Cimatti, L. Hunsberger, A. Micheli, R. Posenato, and M. Roveri. Sound and complete algorithms for checking the dynamic controllability of temporal networks with uncertainty, disjunction and observation. In *21st Intern. Symp. on Temp. Repres. and Reasoning, TIME 2014, Verona, Italy*, pages 27–36, 2014.

[8] C. Combi, L. Hunsberger, and R. Posenato. An algorithm for checking the dynamic controllability of a conditional simple temporal network with uncertainty. In *ICAART 2013 - Proc. of the 5th Intern. Conf. on Agents and Artif. Intell., Vol. 2, Spain, 2013*, pages 144–156, 2013.

[9] Carlo Combi, Mauro Gambini, Sara Migliorini, and Roberto Posenato. Representing business processes through a temporal data-centric workflow modeling language: An application to the management of clinical pathways. *Systems, Man, and Cybernetics: Systems, IEEE Transactions on*, 44(9):1182–1203, Sept. 2014.

[10] Carlo Combi, Matteo Gozzi, Roberto Posenato, and Giuseppe Pozzi. Conceptual modeling of flexible temporal workflows. *ACM Trans. Auton. Adapt. Syst.*, 7(2):19:1–19:29, July 2012.

[11] C. Comin, R. Posenato, and R. Rizzi. A tractable generalization of simple temporal networks and its relation to mean payoff games. In *21th International Symposium on Temporal Representation and Reasoning (TIME 2014)*, Verona, Italy, Sept 2014.

[12] C. Comin, R. Posenato, and R. Rizzi. Hyper temporal networks. *CoRR, (accepted and appearing in [Constraints, an International Journal (2016)])*, abs/1503.03974, 2015.

[13] Carlo Comin and Romeo Rizzi. Dynamic consistency of conditional simple temporal networks via mean payoff games: a singly-exponential time DC-Checking. In *22nd International Symposium on Temporal Representation and Reasoning (TIME 2015)*, pages 19–28. IEEE CPS, sep 2015.

[14] Carlo Comin and Romeo Rizzi. Improved pseudo-polynomial bound for the value problem and optimal strategy synthesis in mean payoff games. *Algorithmica*, pages 1–27, 2016.

[15] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. The MIT Press, 2001.

[16] Rina Dechter, Itay Meiri, and Judea Pearl. Temporal constraint networks. *Artificial Intelligence*, 49(1–3):61–95, 1991.

[17] Johann Eder, Euthimios Panagos, and Michael Rabinovich. Time constraints in workflow systems. In Matthias Jarke and Andreas Oberweis, editors, *Advanced Information Systems Engineering*, volume 1626 of *LNCS*, pages 286–300. Springer Berlin Heidelberg, 1999.

[18] Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *Int. Journal of Game Theory*, 8(2):109–113, 1979.

[19] L. R. Ford and D. R. Fulkerson. *Flows in networks, by L.R. Ford, Jr. [and] D.R. Fulkerson*. Princeton University Press Princeton, N.J, 1962.

[20] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata Logics, and Infinite Games: A Guide to Current Research*. Springer-Verlag New York, Inc., New York, NY, USA, 2002.

[21] L. Hunsberger, R. Posenato, and C. Combi. The dynamic controllability of conditional stns with uncertainty. In *Proc. of the Plan. and Plan Exec. for Real-World Syst.: Princip. and Pract. (PlanEx)*, ICAPS-2012, page 121–128, Atibaia, Sao Paulo, Brazil, 2012.

[22] Luke Hunsberger, Roberto Posenato, and Carlo Combi. A sound-and-complete propagation-based algorithm for checking the dynamic consistency of conditional simple temporal networks. In *22nd International Symposium on Temporal Representation and Reasoning, TIME 2015, Kassel, Germany, September 23-25, 2015*, pages 4–18, 2015.

[23] M. D. Moffitt and M. E. Pollack. Applying local search to disjunctive temporal problems. In *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence, Edinburgh, Scotland, UK*, pages 242–247, 2005.

[24] A. Oddi. Constraint-based strategies for the disjunctive temporal problem: Some new results. In *Proceedings of the Sixth European Conference on Planning*, 2014.

[25] A.K. Pani and G.P. Bhattacharjee. Temporal representation and reasoning in artificial intelligence: A review. *Mathematical and Computer Modelling*, 34(1–2):55–80, 2001.

[26] D.E. Smith, J. Frank, and A.K. Jónsson. Bridging the gap between planning and scheduling. *Knowledge Engineering Review*, 15(1):47–83, 2000.

[27] I. Tsamardinos and M. E. Pollack. Efficient solution techniques for disjunctive temporal reasoning problems. *Artif. Intell.*, 151(1-2):43–89, 2003.

[28] I. Tsamardinos, T. Vidal, and M. Pollack. Ctp: A new constraint-based formalism for conditional, temporal planning. *Constraints*, 8(4):365–388, 2003.

[29] U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158:343–359, 1996.